

THE INFINITE RETINA



Foreword By: Steve Wozniak

**Navigate Spatial Computing, Augmented and Mixed Reality,
and the next tech revolution**

Second Edition

**Irena Cronin
Robert Scoble**

The Infinite Retina

Second Edition

Navigate Spatial Computing, Augmented and Mixed Reality, and the next tech revolution

Irena Cronin

Robert Scoble

⟨packt⟩

The Infinite Retina

Second Edition

Copyright © 2025 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the authors, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

Portfolio Director: Rohit Rajkumar

Relationship Lead: Kaustubh Manglurkar

Program Manager: Sandip Tadge

Content Engineer: Shazeen Iqbal

Technical Editor: Kushal Sharma

Copy Editor: Safis Editing

Indexer: Hemangini Bari

Proofreader: Safis Editing

Production Designer: Pranit Padwal

Growth Lead: Namita Velgekar

Marketing Owner: Nivedita Pandey

First published: April 2020

Second edition: February 2025

Production reference: 1210225

Published by Packt Publishing Ltd.

Grosvenor House

11 St Paul's Square

Birmingham

B3 1RB, UK.

ISBN 978-1-83620-483-1

www.packt.com

Foreword

When I was young, I told my dad that I'd rather own a computer than a home. Of course, that was back when a 4K Data General Nova cost more than a house in Silicon Valley.

Today, just a few decades later, we all carry a computer that is thousands of times more powerful than the computers of my childhood. Computers used to be huge, but today's tech is getting so small, we can wear them. The \$250 Apple AirPods Pros have more compute power inside them than an iPhone 4 and certainly a lot more than that Data General computer that cost tens of thousands of dollars back in the 1970s.

People often ask me what I would engineer if I were starting out today. When I worked at HP we built calculators and huge laser printers. Currently, there's tons of research being done in Spatial Computing, which includes the hardware and software surrounding Robotics, Autonomous Vehicles, Augmented Reality, and Virtual Reality. All of these bring interesting engineering problems that would keep me interested for sure. I see the potential in all of them to bring computing to people in a way that is ubiquitous, more powerful, and far more personal.

Spatial Computing will move computing from something that sits on a desk, or that you hold in your hands, to something you move around with. There, you'll be joined by robots, virtual beings, and assistants. I'm optimistic that Spatial Computing will be used to serve us with a deeper understanding of ourselves, and give us wild new ways to experience the world and new ways to become more productive in businesses and classrooms.

Children soon will have "imagination machines" on their faces. That will be in contrast to the toys I grew up with. One of my favorites was a transistor radio, which had six transistors (a modern phone has billions of transistors, by comparison). That radio brought the world to me. Spatial Computing will open up new worlds to today's kids that simply didn't exist when I was growing up.

To me I'm excited by what the children of today will do with these technologies and hope they remember that life is about happiness. Smiling and laughing is good. Hearing music in new ways is good. Telling new jokes to each other is good. Spatial Computing can do all that and can bring new ways to learn, work together, and travel the world to boot.

What Irena Cronin and Robert Scoble have done in their book *The Infinite Retina* is provide readers with an accessible, useful and entertaining guide to all of the latest advances in Spatial Computing and how this new technology is influencing and priming the way forward for a host of new innovations that will affect how we live our lives, how we go to work, and how we communicate with each other for years to come.

- *Steve Wozniak*

Contributors

About the authors

Irena Cronin is SVP of Product for DADOS Technology which is creating a data analytics and visualization platform for the Apple Vision Pro. She is also the CEO of Infinite Retina which provides research to help companies develop and implement AI, AR and other new technologies for their businesses.

Cronin has a joint MBA/MA from the University of Southern California and an MS with Distinction in Management and Systems from New York University. She graduated with a BA from the University of Pennsylvania with a major in Economics (summa cum laude).

In the memory of my husband, Danny, who loved all things tech.

Robert Scoble is a renowned technology futurist, author, and industry thought leader with a deep focus on artificial intelligence, augmented reality, and emerging technologies. Best known for his influential blogging and early-stage tech reporting in Silicon Valley, he has played a key role in introducing groundbreaking startups to the world.

As a former evangelist at Microsoft and Rackspace, Scoble has worked closely with some of the most innovative minds in the industry, providing unique insights into the evolution of AI, spatial computing, and digital transformation.

Thank you to all the engineers and pioneers who have been working on these things, in some cases, for decades.

About the reviewer

Alessandro “Alex” Straccia is a Creative Technologist with a background in Advertising, Marketing, and Computer Science. With over 25 years of experience, he has worked on digital projects, as well as augmented reality (AR) and virtual reality (VR) projects for major brands and studios. As a certified startup board member and advisor, he helps guide early-stage companies in Creative Tech. Alessandro is also the author of Interactive Web Development with Three.js and A-Frame and lectures on immersive technologies.

I would like to thank my mom and dad who taught me how to be the person I turned into. And to my kids Victor and Lucca and my wife Aline who teaches me how to be the person I’m becoming.

Table of Contents

Introduction	xxiii
<hr/>	
Part I: Why Spatial Computing and Why Now?	1
<hr/>	
Chapter 1: Prime Directive	3
<hr/>	
What Makes Us Human?	5
Drivers and Benefits	7
Potential Dangers	8
Understanding the Natural World	10
Faster Wireless Leads to New Affordances	12
Data, Data, Data Everywhere	13
Electric Sheep Robotics Dreams	15
Dancing into Different Worlds	17
The Frustrated Pioneer	19
Breakthroughs in Sensors	21
Driven to Improve	24
Hands-On Use	27
Beginning	29
References	30

Chapter 2: Four Paradigms and Six Technologies	31
The Four Paradigms	31
Paradigm One—The Arrival of the Personal Computer • 33	
Paradigm Two—Graphical Interfaces and Thinking • 36	
Paradigm Three—Mobile • 38	
Paradigm Four—Spatial Computing • 40	
Six Technologies	42
Optics and Displays • 42	
Wireless and Communications • 46	
Control Mechanisms (Voice, Eyes, and Hands) • 48	
Sensors and Mapping • 51	
Computing Architectures • 56	
AI (Decision Systems) • 59	
Evolution	62
References	62
Chapter 3: The Impact of Generative AI on Spatial Computing	63
Enhancements in Immersive Experiences	63
Hyper-Realistic Simulations • 63	
Emotional Engagement and Behavioral Training • 65	
Customization and Personalization	66
Interior Design and Architecture • 66	
Education and Training • 67	
Gaming • 69	
Beyond Gaming • 72	
Retail and Shopping • 74	
Personalized Work Environments • 74	
Future Directions for Customization and Personalization • 76	
Dynamic Content Generation in More Detail	76
Adaptive AR Applications • 78	
Real-Time Data Visualization • 80	

Accessibility	82
Data Synthesis and Simulation	84
Challenges and Ethical Considerations	85
Data Quality and Bias • 85	
Ethical Considerations • 87	
Educating Stakeholders • 89	
Future Directions	89
References	90

Part II: The Seven Visions 91

Chapter 4: Vision One: Transportation Automates	93
A Future Vision	93
The Paths to Autonomy	95
Falling Car Transportation Costs	99
Autonomous Vehicle Round-up	101
Zoox • 102	
Ford • 105	
General Motors' Cruise • 106	
Waymo • 107	
Tesla • 109	
Apple • 110	
Hyundai, Pony.ai, Aptiv, and Aurora • 110	
Autonomous Trucks • 111	
Flying Cars (eVTOLs) • 113	
The Autonomous and Electric Culture Shock	115
The Future City	121
Auto Dealers • 122	
Smart Roads • 122	
Interstates • 123	

Factories • 123	
Real Estate • 124	
Remote Working • 124	
Travel Times and Ranges • 124	
Navigation • 125	
Green Thinking • 125	
AI Assistants • 125	
Entertainment • 126	
Fewer Cars • 126	
“Data Bubbles” Will Abound	126
Chapter 5: Vision Two: Virtual Worlds Appear	129
From 2D to 3D	129
Increasing the Bandwidth • 130	
Change is Coming • 131	
Emergence as a Human Hook	133
The Brain and VR • 133	
The Right Tool for the Job • 134	
The “Glassholes” Show the Way	134
Defeated Expectations • 134	
Where Did Things Go Wrong? • 136	
What It Means to Be Human • 137	
Embodiment, Immersion, Presence, Empathy, and Emergence	139
Embodiment and Immersion • 139	
Evolution • 140	
Meta Has the Edge • 141	
Market Evolution • 142	
2014’s New Wave • 144	
Inside-Out • 147	
Mistakes Were Made • 148	

The Spectrum: Augmented Reality from Mobile to Mixed Reality	150
Closer Computers • 150	
The Origins and Limitations of Augmented Reality • 151	
Wearable AR • 154	
Too Much and Too Little • 157	
Getting Past the Setbacks • 158	
Vergence and Accommodation • 159	
Fast, Cheap, or Good • 160	
Differences to Be Aware of • 163	
From Kinect to the HoloLens • 164	
Technology Takes a Leap • 166	
Patents and Problems • 167	
Apple Vision Pro • 168	
Volumetric and Light Field – Capturing Our World in a New Way	169
Shrinking Concepts, Expanding Horizons • 169	
Shooting Some Volumetric • 170	
Camera Magic • 172	
AR Clouds (3D Mapping) and Internet Infrastructure Including 5G	174
A New Kind of Map • 174	
Building a Map • 175	
3D Gaussian Splats • 176	
<i>Understanding 3D Data • 176</i>	
<i>What Does Gaussian Mean? • 176</i>	
<i>The Concept of Splats • 177</i>	
<i>Combining the Concepts • 177</i>	
Ready for 5G • 178	
New Audio Capabilities Arrive	179
Audio Augmentation • 179	
Spatial Audio • 180	

Gaming and Social VR Experiences	181
A World Without Rules • 181	
The Ripples of the VR Phenomenon • 183	
Building a VR Community • 184	
Cinematic and Interactive Experiences	185
Festivals That Include VR • 185	
Content Production and Previsualization	186
The Merging of the Virtual and the Real	187
References	187
Chapter 6: Vision Three: Augmented Manufacturing	189
Robots and Other Augmentation	190
Augmented Reality Visions and Informed Robots	196
The Mirror World Comes to Factory Floors	204
The Human Robot	207
References	208
Chapter 7: Vision Four: Robot Consumers	209
What Drives Consumers?	209
Getting Emotional • 210	
Now in 3D • 210	
Retail Destruction? No, But Radical Change Underway	212
Trying to Keep Up • 213	
Looking Ahead • 214	
Variations on a Retail Theme • 215	
<i>The Integration of Spatial Computing in Retail—Transforming Customer Experiences, Operations, and Business Strategies</i> • 217	
Spatial Computing for Logistics	224
Spatial Computing Delivery Methods—Robots, Drones, Future • 226	
In Your Living Room and Elsewhere—Spatial Computing Marketing Experiences	231
Starting in the Living Room • 231	

A Personalized Experience • 232	
The VR World • 233	
Making an Impact • 235	
A “No Install Needed” AR World • 238	
The Past and The Future	241
Will Anything Remain the Same?	246
References	246
 Chapter 8: Vision Five: Virtual Healthcare 247	
The Virtualist	248
Fixing Humans—Surgery Preparation and Surgery	251
Focus on MediView • 254	
<i>FDA Clearance and Regulatory Milestones</i> • 254	
<i>Partnership with GE</i> • 255	
<i>XR Products and Technological Advancements</i> • 255	
<i>Expanding into Surgical Navigation and Cardiology</i> • 256	
<i>AI Integration and Future Directions</i> • 256	
Brain Tricks/Virtual Medicine	257
The Virtual Coach and the “Check Engine Light”	263
Data, Data, Everywhere	265
Precision and Personalization	268
References	268
 Chapter 9: Vision Six: Virtual Trading and Banking 269	
3D Data Visualization	269
Virtual Trading	272
ATM Security and Facial Payment Machines	275
Virtual Branch Functionality and Customer Service	281
Future Adoption	283
References	283

Chapter 10: Vision Seven: Real-Time Learning	285
Experiential Digital Learning	286
No Teachers, Books, or Movies? New Approaches in the Offing • 288	
Brainier Learning	289
Spotlight on Cornerstone • 292	
The Future of Education	295
The Grabbable Classroom • 296	
Workplace Training and Simulation • 301	
A Virtual Arm for Human Resources	309
Our Learning Futures	310
References	311
Part III: The Spatial Business	313
Chapter 11: The Always Predicted World	315
The Predictive Breakthrough	315
Glasses That Understand You • 316	
(Human) Assistance Not Required • 317	
<i>Pervasive</i> • 319	
<i>Proactive</i> • 321	
<i>Connected</i> • 324	
Data Dance with Seven Verticals	326
Transportation • 327	
Technology, Media, and Telecommunications (TMT) • 330	
Manufacturing • 332	
Retail • 333	
Healthcare • 335	
Finance • 336	
Education • 336	

The Bubbleland	339
The Bubble Network • 339	
Unified Learning, Shared Experiences • 339	
Pervasive and Proactive Future Assistance	340
Chapter 12: The Business of Generative AI	341
Generative AI in General	342
Content Creation and Media Automation • 342	
Product Design and Prototyping • 343	
AI-Driven Customer Interactions • 343	
Expanding Creativity in Entertainment and the Arts • 344	
Revolutionizing Education and Training • 345	
Healthcare Innovation and Medical Research • 346	
The Convergence of Generative AI and Spatial Computing	346
Automated Content Generation for Immersive Environments • 347	
<i>Game Development: Creating Dynamic Worlds and Realistic Characters</i> • 347	
<i>Film and Entertainment: Virtual Sets and Digital Doubles</i> • 348	
<i>Architecture and Real Estate: Immersive Visualizations and Virtual Tours</i> • 349	
<i>Education and Training: Interactive Learning Environments</i> • 350	
<i>Marketing and Retail: Dynamic Product Showcases and AR Experiences</i> • 351	
<i>Automotive Industry: Automated Prototyping and Virtual Testing</i> • 351	
Personalized Customer Experiences through Data-Driven Insights • 352	
<i>AI's Comprehensive Role in Understanding Customer Behavior</i> • 353	
<i>Spatial Computing: Bridging Insights and Real-Time Experiences</i> • 353	
<i>Revolutionizing Entertainment with AI-Driven Spatial Computing</i> • 354	
<i>Creating Competitive Advantage through Immersive, Personalized Experiences</i> • 354	
Smart Cities and Urban Development: Real-Time Digital Twins • 356	
<i>Understanding Digital Twins: A Complex Synthesis of Data and Spatial Computing</i> • 356	
<i>Generative AI: Transforming Digital Twins into Intelligent Urban Systems</i> • 356	
<i>Optimizing Traffic Flow and Reducing Congestion</i> • 357	
<i>Enhancing Utility Management and Infrastructure Maintenance</i> • 357	

<i>Business Opportunities in the Digital Twin Ecosystem</i> • 357	
<i>Predictive Maintenance and Infrastructure Optimization</i> • 358	
<i>Urban Planning and Development Consultancies</i> • 358	
<i>Transforming Public Safety and Emergency Response</i> • 358	
<i>The Ethical Considerations and Privacy Challenges of Smart Cities</i> • 359	
<i>Looking Ahead: The Future of Smart Cities and Digital Twin Technology</i> • 359	
Economic Impact	360
Automation and Cost Efficiency • 360	
Enhanced Productivity and Streamlined Operations • 361	
Emergence of New Business Models	361
Digital Assets and Virtual Economies • 362	
Economic Impact on Traditional Industries • 362	
Long-Term Economic Trends and Future Projections • 363	
Challenges and Considerations	364
Ethical and Privacy Concerns • 364	
High Costs and Accessibility Barriers • 365	
Technical Complexity and Integration Challenges • 365	
Intellectual Property and Copyright Issues • 366	
<i>The Ambiguity of Ownership: Who Owns AI-Generated Content?</i> • 366	
<i>Legal and Operational Risks for Businesses Utilizing Generative AI</i> • 367	
<i>Authenticity, Traceability, and the Role of Blockchain Technology</i> • 367	
<i>Establishing a Legal Framework for AI-Generated Content</i> • 368	
<i>The Path Forward: Balancing Innovation with Legal Certainty</i> • 369	
Regulatory Compliance and Global Standards • 369	
<i>The Regional Disparities in Regulatory Approaches</i> • 369	
<i>The Rise of AI Ethics Regulations and the Challenge of Responsible AI</i> • 370	
<i>The Role of Compliance Programs and Real-Time Monitoring</i> • 371	
<i>The Consequences of Non-Compliance: Financial and Reputational Risks</i> • 372	
<i>Proactive Engagement with Regulatory Bodies and Evolving Legislation</i> • 372	
<i>Investing in RegTech Solutions: Automating Compliance for Efficiency</i> • 373	
<i>The Future of Compliance: Toward Harmonization and Global Standards</i> • 373	

Environmental Impact and Sustainability • 374	
<i>The Environmental Costs of AI and Spatial Computing</i> • 374	
<i>Mitigating the Environmental Impact: Strategies and Solutions</i> • 375	
<i>Incorporating Sustainable Practices into Business Models</i> • 376	
<i>Leveraging Green Technology as a Competitive Advantage</i> • 376	
<i>The Role of Policy and Industry Collaboration in Promoting Sustainability</i> • 377	
<i>The Future of AI and Spatial Computing in an Eco-Conscious World</i> • 378	
Workforce Adaptation and Skill Gaps • 378	
The Future of Generative AI and Spatial Computing in Business	379
Chapter 13: The Impact of the Apple Vision Pro	381
Technological Features	382
<i>VisionOS and User Interface Innovations</i> • 382	
<i>Eye-Tracking, Gesture Control, and Passthrough Capabilities</i> • 382	
<i>Micro-OLED Displays and Hardware Advancements</i> • 383	
<i>Apple's Vision for Immersive Technology</i> • 383	
<i>Reimagining Computing Paradigms</i> • 383	
<i>Developer Empowerment</i> • 383	
<i>Ecosystem Integration</i> • 384	
<i>Inspiring a New Era of Content Creation</i> • 384	
Redefining Interaction in Spatial Computing	384
Natural User Interfaces • 384	
Eye-Tracking: Precision and Context Awareness • 385	
Gesture Recognition: Intuitive Interaction • 385	
<i>Voice Input: Conversational Control</i> • 385	
Transition from Screens to Spaces • 386	
From Flat Screens to 3D Interfaces • 386	
Passthrough and MR • 386	
Collaboration in Shared Spaces • 387	
UX Design • 387	
Intuitive Design Principles • 387	

Accessibility Features • 387	
Reducing Cognitive Load • 388	
Advancing UX Design in Spatial Computing • 388	
Applications of the Vision Pro across Industries	388
Entertainment and Media • 389	
<i>Immersive Streaming and 3D Cinema Experiences • 389</i>	
<i>Gaming in Augmented and Mixed Reality • 389</i>	
Productivity and Collaboration • 389	
<i>Virtual Workspaces and Remote Collaboration Tools • 389</i>	
<i>Enhanced Presentations and Prototyping in 3D Spaces • 390</i>	
Healthcare • 390	
<i>Training and Simulations for Medical Professionals • 390</i>	
<i>Patient Care and Rehabilitation through Immersive Experiences • 391</i>	
Education • 391	
<i>Transforming Classrooms and Virtual Learning with Interactive 3D Content • 391</i>	
<i>Democratizing Access to Immersive Educational Tools • 391</i>	
Retail and E-Commerce • 392	
<i>AR-Powered Shopping and Virtual Product Experiences • 392</i>	
<i>Enhancing Customer Engagement through Spatial Computing • 392</i>	
Transforming Content Creation and Development with the Apple Vision Pro	393
Creative Possibilities: Redefining Opportunities for Filmmakers, Designers, and Developers • 393	
For Filmmakers: Immersive Storytelling • 393	
For Designers: Enhanced 3D Modeling and Visualization • 393	
For Developers: Pioneering Interactive Applications • 394	
Spatial App Development: Empowering Developers with VisionOS and Apple's Tools • 394	
Developer Tools and SDKs • 394	
Opportunities for Developers • 395	
3D Content Ecosystem: Promoting Mass Adoption of 3D Content Creation • 395	
Lowering Barriers to Entry • 395	
Fostering Collaboration • 396	

Driving Consumer Adoption • 396	
Economic and Market Impact of the Apple Vision Pro	396
Consumer Adoption: Driving Spatial Computing into the Mainstream • 397	
Appealing to Early Adopters and High-End Consumers • 397	
Lowering Barriers to Entry over Time • 397	
Mainstream Appeal through Ecosystem Integration • 397	
Impact on Competitors: Shaping the Spatial Computing Market • 398	
<i>Setting a New Benchmark for Quality</i> • 398	
<i>Driving Innovation across the Industry</i> • 398	
<i>Expanding the Market</i> • 398	
Market Opportunities: Unlocking New Business Models and Ecosystems • 399	
<i>Monetizing Spatial Applications</i> • 399	
<i>Fostering Collaboration and Innovation</i> • 399	
<i>Accelerating Content Creation</i> • 400	
Challenges and Criticisms of the Apple Vision Pro	400
Technical Limitations • 400	
<i>Battery Life</i> • 400	
<i>Weight and Comfort</i> • 401	
<i>Processing Power and Heat Management</i> • 401	
<i>Adoption Barriers</i> • 401	
<i>Cost Concerns</i> • 401	
<i>Accessibility for Diverse Demographics</i> • 402	
Privacy and Ethical Considerations • 402	
<i>Risks Associated with Eye-Tracking and Data Collection</i> • 402	
<i>Ethical Implications of Immersive Technology on User Well-Being</i> • 402	
<i>Balancing Innovation and Responsibility</i> • 403	
Vision Pro and the Future of Spatial Computing	403
Advancing Mixed Reality • 404	
<i>Blurring the Lines between AR and VR</i> • 404	
<i>Setting a New Benchmark for Hardware</i> • 404	
<i>Inspiring Future MR Devices</i> • 404	

Integration with Emerging Technologies • 405	
<i>Artificial Intelligence and Machine Learning</i> • 405	
<i>Integration of IoT</i> • 405	
<i>Expanding the Vision Pro's Capabilities</i> • 405	
Impact on the Spatial Computing Ecosystem • 406	
<i>Growth of Interconnected 3D Environments</i> • 406	
<i>Vision Pro as a Foundational Step toward the "Spatial Web"</i> • 406	
<i>Accelerating Ecosystem Growth</i> • 406	
Case Studies and Early Use Cases of the Apple Vision Pro	407
Developer Insights: Building the Future of Spatial Applications • 407	
<i>Early Application Development</i> • 407	
<i>Challenges and Opportunities for Developers</i> • 408	
<i>Consumer Reactions: Early Feedback and User Experiences</i> • 408	
<i>Positive Reception</i> • 408	
<i>Concerns and Criticisms</i> • 409	
<i>Anticipation for Future Updates</i> • 409	
Industry Adoption: Pilot Projects and Early Use Cases • 409	
<i>Healthcare: Transforming Training and Patient Care</i> • 409	
<i>Education: Immersive Learning Experiences</i> • 410	
<i>Entertainment: Redefining Media Consumption</i> • 410	
<i>Corporate Environments: Productivity and Collaboration</i> • 410	
The Future of the Vision Pro	411
Summarizing Vision Pro's Impact • 411	
Looking Ahead: The Next Decade of Immersive Technology • 411	
<i>Paving the Way for Mainstream Adoption</i> • 412	
<i>Revolutionizing Work and Communication</i> • 412	
Call to Action for Developers and Innovators • 412	
<i>Embrace the VisionOS Ecosystem</i> • 412	
<i>Explore Diverse Use Cases</i> • 413	
<i>Advocate for Ethical Development</i> • 413	

Chapter 14: How Human?	415
Transformative Experience	416
Privacy and Security	416
On Facebook and Cambridge Analytica • 416	
OJ Simpson and Location Data • 421	
Apple and the FBI • 422	
Identity and Ownership	423
Blurred Faces on Google Street View • 423	
China and Facial Recognition • 424	
Issues with Facial Recognition in the US and Europe • 426	
Human Social Good	430
Other Books You May Enjoy	435
Index	437

Introduction

Spatial Computing comprises all software and hardware technologies that enable humans, virtual beings, or robots to move through real or virtual worlds, and includes Artificial Intelligence, Computer Vision, Augmented Reality (AR), VR, Sensor Technology, and Automated Vehicles.

Seven industry verticals will see transformational change due to Spatial Computing: Transportation; Technology, Media, and Telecommunications (TMT); Manufacturing; Retail; Healthcare; Finance; and Education.

These changes are what is driving strategy at many tech companies and the spending of billions of dollars of R&D investment. Already, products such as Microsoft's HoloLens AR headset have seen adoption in places from surgery rooms to military battlefields. Devices like this show this new computing paradigm, albeit in a package that's currently a little too bulky and expensive for more than a few of the hardiest early adopters. However, these early devices are what got us to be most excited by a future that will be here soon.

Our first experiences with HoloLens, and other devices like it, showed us such a fantastic world that we can predict that when this new world arrives in full, it will be far more important to human beings than the iPhone was.

We were shown virtual giant monsters crawling on skyscrapers by Metaio years ago near its Munich headquarters. As we stood in the snow, aiming a webcam tethered to a laptop at the building next door, the real building came alive thanks to radical new technology. What we saw at Metaio had a similar effect on Apple's CEO, Tim Cook. Soon, Apple had acquired Metaio and started down a path of developing Augmented Reality and including new sensors in its products, and the capabilities of this are just starting to be explored. Today's phones have cameras, processors, small 3D sensors, and connectivity far better than that early prototype had, and tomorrow's phones and soon, the glasses we wear will make today's phones seem similarly quaint.

In Israel, we saw new autonomous drones flying over the headquarters of Airobotics. These drones were designed to have no human hands touching them. Robots even changed memory cards and batteries. New Spatial Computing technology enabled both to “see” each other and sense the world around us in new ways. The drones were designed to fly along oil pipelines looking for problems, and others could fly around facilities that needed to be watched. Flying along fences and around parking lots, their Artificial Intelligence could identify things that would present security or other risks. These drones fly day or night and never complain or call in sick.

Focusing on just the technology, though, would have us miss what really is going to happen to the world because of these technologies. Our cities and countryside will reconfigure due to automation in transportation and supply chains as robot tech drives our cars, trucks, and robots rolling down sidewalks delivering products. We’ll spend more time in virtual worlds and metaverses. More of our interfaces, whether they are the knobs on our watches, cars, doors, and other devices, will increasingly be virtualized. In fact, many things that used to be physical may be virtualized, including stores and educational lessons from chemistry experiments to dissection labs.

Computing will be everywhere, always listening, always ready to talk back, and once we start wearing Spatial Computing glasses, visual computing will always be ready to show us visualizations of everything from your new designs to human patterns in stores, on streets, and in factories and offices. Some call this “invisible computing” or “ambient computing,” but to us, these systems that use your eyes, voice, hands, and even your body as a “controller” are part of Spatial Computing.

At the same time, all this new computing is joined by radically fast new wireless technology in the form of 5G. The promises of 5G are threefold. First, we’ll have more than a gigabit of data if we have the highest bitrates, and even the lowest rates promise to give us more bandwidth than current LTE phones. Second, wireless will soon add almost no latency, which means that as soon as you do something, like throw a football in a virtual game, it’ll happen, even if you have many players viewing the football in real time. Third, 5G supports many more devices per tower, which means you will be able to live stream even a Taylor Swift concert with tens of thousands of fans filling a stadium.

When you combine 5G with all the new things under the Spatial Computing umbrella, you get a big bang. All of a sudden, cars can work together, where one of the cars will be sending detailed 3D imaging of streets instantly to other cars behind it. New kinds of virtual games will be possible in the streets, where hundreds of people can play virtual football games in parks and other places. Crazy new virtual shopping malls appear where virtual celebrities show you around, and you can imagine new products directly in a virtual scan of your home and in other places as well.

A range of new capabilities will appear over the next few years in devices you wear over your eyes. There will be very light ones that are optimized for utility, showing you navigation, notifications, reminding you where you have left things, or nagging you to do some exercise or meditation to keep on top of your physical and mental health. There also will be heavier devices that will be more optimized for everything from detailed design or architecture work to entertainment and video game work. We can even imagine owning several different kinds of Spatial Computing devices, along with some smart contact lenses, that will let us go out on a date night without looking like we have any computing devices on at all.

We currently have VR devices that cost a few hundred dollars that are great for games and a few other things, like corporate training. On the more expensive side of the scale, we have devices that can be used by car designers, or even as flight simulators to train airline pilots. The expensive ones, though, will soon look as out of date as one of the first cell phones does today. By 2028, the computing inside will shrink to a fraction of the size of today's devices and the screens inside will be much sharper and capable of presenting virtual and augmented worlds to us that far exceed what we can experience today.

It is this next wave of devices that will usher in the paradigm shift in computing and in human living that we are discussing here. Already these changes are benefiting many enterprises, raising productivity. Inside many warehouses, hundreds of thousands of robots scurry about, moving products from trucks to packages. These new warehouses have evolved over the past decade and enable retailers to keep up with floods of new online orders that, back in 2000, were only dreamt of in futuristic books like this one.

The productivity gains will spread to many jobs. At Cleveland Clinic, surgeons are already using similar technology that shows them digital views from ultrasound, CAT scans, and other sensors. Like the warehouse worker who sees a blue line on the floor telling her how to find the product she's looking for, in this case, when a surgeon navigates to the right place to cut out a cancerous tumor, it lights up like a missile guidance system and tells the surgeon they are in the right place.

Other systems help workers “phone a friend” with new remote assistance features. This can help companies that have expensive machinery, or other work forces, including surgeons, architects, and engineers, save money. At some plants, the savings will be substantial. It took us 30 minutes to simply walk across the Boeing floor where it builds airliners. Asking someone for advice virtually might save someone an hour of walking just to come over and see your problem in a plant like that.

New devices let these remote helpers see what you are dealing with, and they can often show you visually what to do. Imagine trying to remove an engine while holding a phone or tablet in your hand. These systems, because they use wearable glasses, can let workers use both of their hands while talking and showing the remote assistant what is happening. The savings in downtime can be extreme. Imagine a problem that is causing a shutdown in a line at Ford. Every minute it is down, Ford loses about \$50,000.

Even for salespeople and managers, the cost savings add up. A flight from San Francisco to Los Angeles usually costs about \$700, including airfare, hotel, a decent meal, and an Uber/taxi ride or two. Increasingly, these meetings will be replaced by ones held in Virtual Reality. Two headsets are less than \$800. In just a couple of virtual meetings, the headsets pay for themselves. Social VR from Facebook, Spatial, Microsoft, and others, are rapidly improving to make these virtual meetings almost as good as physical ones. The time and cost saved will add up to big numbers, and workers will be happier. Workforces will be less likely to pick up viruses from travelers too, which will also add up to big savings for corporations and reduced risks.

More lives could be saved, too. Mercedes-Benz had an Augmented Reality system built to show first responders how to cut apart a wrecked car. The app showed where fuel and electrical lines were so that firefighters working to free an accident victim wouldn't start a fire or electrocute themselves. This isn't the only example we have of this technology helping first responders, either through better training or by giving them various assistance on scene. One such system helps police gather evidence and then be able to recreate crime scenes for juries that they can virtually walk around.

Here, we've given you a taste of just how much the world is about to get reconfigured because of Spatial Computing technologies. Let's dig into what this book will cover.

Who this book is for

We wrote *The Infinite Retina* for a wide audience of non-technical people, rather than for engineers. By the end of the book, you should understand the technologies, companies, and people who are changing computing from being something you do while sitting and facing a flat computer screen or while holding a phone, to computing that you could move through three-dimensionally. We focus on how Spatial Computing could be used by enterprises and effectively radically change the way human beings learn from information and visuals and understand their world. We very strongly believe that enterprise use of Spatial Computing will lead to massive consumer use, and we are excited to share our learning in this book with you.

What this book covers

Part I: Why Spatial Computing and Why Now?, The Significance of Spatial Computing

Chapter 1, Prime Directive, Mobile phones soon will give way to headsets and glasses that bring computing to every surface. What is driving all of this new technology? We have a need for complex technologies to keep us around on this planet longer and in a more satisfied and productive state. What will drive us to build or buy new headsets, sensors, and vehicles, along with the connected systems controlled by Artificial Intelligence? Augmentation is coming, and that can mean a lot of different things, which we will explore.

Chapter 2, Four Paradigms and Six Technologies, we will look at the previous three foundations of personal computing and include the new Spatial Computing paradigm. The six technologies discussed are those that enable Spatial Computing to work: Optics and Displays, Wireless and Communications, Control Mechanisms (Voice and Hands), Sensors and Mapping, Compute Architectures (new kinds of Cloud Computing, for instance), and Artificial Intelligence (Decision Systems).

Chapter 3, The Impact of Generative AI on Spatial Computing, we discuss how Generative AI is revolutionizing Spatial Computing by enhancing immersive experiences, enabling hyper-realistic simulations, and personalizing interactions across industries. Generative AI plays a pivotal role in training and education by creating highly detailed virtual environments for aviation, medical, and military applications. AI-driven simulations now incorporate emotional and psychological realism, making training more effective and preparing professionals for real-world scenarios.

Part II: The Seven Visions, A Road Map to Spatial Computing Disruption in Seven Industries

Chapter 4, Vision One – Transportation Automates, where you will discover that soon you will tell your glasses, “Hey, I need a ride,” and you’ll see your ride arrive. Sounds like Uber or Lyft, right? Look closer, there isn’t a driver inside. Now think of the cost and the other advantages of that. Economists see that such a system could be a fraction of the cost, and could do many other things as well: “Hey car, can you go pick up my laundry and then dinner for our family?” The problem with such a world is that it is probable, many tell us, that we’ll see much more traffic near cities as we use transportation to do new things, like pick up our laundry. This is why Elon Musk came up with the Boring Company to build tunnels under cities. We show some other solutions pioneers have come up with, including special roads for these vehicles and new kinds of flying vehicles that will whisk commuters into city centers, passing above all that new traffic.

Chapter 5, Vision Two – Virtual Worlds Appear, provides details on Technology, Media, and Telecommunications, another of our seven industry verticals to be disrupted. We start out by detailing the different kinds of devices that are available to bring a spectrum of Spatial Computing capabilities to your face, from Virtual and Augmented Reality headsets to lightweight smart information glasses, and even contact lenses with displays so small that it will be very hard to tell that your friend is wearing one.

Chapter 6, Vision Three – Augmented Manufacturing, you'll learn about how Spatial Computing is changing how factories are even designed. Increasingly, these factory floors are using robots. The robots are different than they used to be, too. The older ones used to be kept in cages designed to keep humans away. Those can still be found welding, or like in Ford's Detroit factory, putting windshields into trucks. Newer robots work outside cages and sometimes, can even touch humans. These types of robots are called "cobots" because they cohabit with humans and can greatly assist workers.

Chapter 7, Vision Four – Robot Consumers, we detail changes that make retail stores, even traditional ones, more efficient and better for both consumers and sales for producers, and useful new Augmented Reality technologies that make shopping at home much easier.

Chapter 8, Vision Five – Virtual Healthcare, we discuss how the virtual healthcare system guides the surgeon to the right spot to cut out a patient's cancerous tumor.

Chapter 9, Vision Six – Virtual Trading and Banking, we cover the future uses for Spatial Computing in the financial industry.

Chapter 10, Vision Seven – Real-Time Learning, we talk with educators and others who are using technology aggressively to make learning more virtual. It isn't only for kids, either. Soon, because of automation, we'll need to retrain millions of adults around the world, and schools and universities are responding with new curricula, new learning programs for Virtual and Augmented Reality, and new support systems to enable even truck drivers to change careers. Speaking of careers, already at companies like Caterpillar, it is using Augmented Reality glasses to train workers to fix their expensive tractors in real time. Many new VR-based training systems are being developed, from simulators to help police learn how to deal with terrorist situations to ones that show quarterbacks how to perform better, to training at Walmart that shows retail workers how to manage stores better. Verizon even trained its retail store workers on what to do if they are being robbed using VR-based training. What if, though, the system could do even more, we asked, and predict what we might do next and assist us with that?

Part III: The Spatial Business, Productivity, Efficiency, and Making Profits Are About to Undergo Radical Changes

Chapter 11, The Always Predicted World, we show how that data will be used in each of our seven disruptable industries to serve users in radically new ways.

Chapter 12, The Business of Generative AI, we discuss how Generative AI and Spatial Computing are transforming businesses by enhancing creativity, efficiency, and customer engagement. Companies use these technologies to develop innovative products, streamline operations, and gain a competitive edge.

Chapter 13, The Impact of the Apple Vision Pro, we explore the Vision Pro's technological features, applications across industries, content development, market impact, challenges, and its role in the future of Spatial Computing. visionOS, the core of this device, introduces a three-dimensional interface that adapts to user behavior, making digital interactions feel more natural and efficient.

Chapter 14, How Human?, we provide a philosophical framework put forth by L.A. Paul of Yale University in her book, *Transformative Experience*, that explains why human beings tend to have cognitive issues with radical new technologies. We then discuss recent issues regarding privacy, security, identity, and ownership, and how they relate to Spatial Computing. Finally, we take up how Spatial Computing technologies can be utilized to bring about human social good.

To get the most out of this book

Anyone who has an interest in Spatial Computing can read this book. No technical knowledge is assumed.

Download the color images

We also provide a PDF file that has color images of the screenshots/diagrams used in this book. You can download it here: <https://packt.link/gbp/9781836204831>.

Conventions used

There are a number of text conventions used throughout this book.

Bold: Indicates a new term, an important word, or words that you see on the screen. For instance, words in menus or dialog boxes appear in the text like this. For example: “New realities are appearing thanks to **Virtual Reality (VR)** and **Augmented Reality (AR)**.”



Warnings or important notes appear like this.



Tips and tricks appear like this.

Get in touch

Feedback from our readers is always welcome.

General feedback: Email feedback@packtpub.com and mention the book's title in the subject of your message. If you have questions about any aspect of this book, please email us at questions@packtpub.com.

Errata: Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you reported this to us. Please visit <http://www.packtpub.com/submit-errata>, click **Submit Errata**, and fill in the form.

Piracy: If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at copyright@packtpub.com with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit <http://authors.packtpub.com/>.

Share your thoughts

Once you've read *The Infinite Retina, Second Edition*, we'd love to hear your thoughts! Please [click here](#) to go straight to the Amazon review page for this book and share your feedback.

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Download a free PDF copy of this book

Thanks for purchasing this book!

Do you like to read on the go but are unable to carry your print books everywhere?

Is your eBook purchase not compatible with the device of your choice?

Don't worry, now with every Packt book you get a DRM-free PDF version of that book at no cost.

Read anywhere, any place, on any device. Search, copy, and paste code from your favorite technical books directly into your application.

The perks don't stop there, you can get exclusive access to discounts, newsletters, and great free content in your inbox daily.

Follow these simple steps to get the benefits:

1. Scan the QR code or visit the link below:



<https://packt.link/free-ebook/9781836204831>

2. Submit your proof of purchase.
3. That's it! We'll send your free PDF and other benefits to your email directly.

Part I

Why Spatial Computing and Why Now?

The Significance of Spatial Computing

This section comprises the following chapters:

- *Chapter 1, Prime Directive*
- *Chapter 2, Four Paradigms and Six Technologies*
- *Chapter 3, The Impact of Generative AI on Spatial Computing*

1

Prime Directive

A perfect storm of change started to arrive during the early 2020s. We predict that we will see computing switch from staring into a four-inch piece of glass in our hands to using devices on our faces that bring computing to every surface. Along with that, we will see advances such as vehicles moving without humans in them for the first time. A fourth paradigm of the personal computing age is upon us, Spatial Computing, and it is one that truly makes personal computers even more personal.

With the coming of Spatial Computing and its two purest members, **Virtual Reality (VR)** and **Augmented Reality (AR)**, fed by **artificial intelligence (AI)**, we believe that businesses and human cooperation need to be aimed at one thing: working together to build complex technologies to keep us around on this planet longer and in a more satisfied and productive state, while paying attention to the effects that these technologies have on ourselves and the planet. Spatial Computing in the 2020s will see immense challenges, but great opportunities will be available for brands to use new technologies for combined social good. Just what is it about Spatial Computing that is making human beings crave it?

At stadiums around the world, they are getting ready for it. We visited a place in Las Vegas because of it. Hollywood is getting ready for it. We visited a huge studio in Manhattan Beach, California, to see it. Investors and entrepreneurs are getting ready for it. We visited several in Seattle, New York, Los Angeles, and Silicon Valley. Hospitals, shopping malls, hotels, cities, automotive companies, power companies, banks, and more are investing in it. Tens of billions of dollars are being invested in it by both big and small tech companies.

That “it,” which is Spatial Computing, is computing that you, virtual beings, or robots can move around in. It includes all the software and technology needed to move around in a digital 3D world. It’s very different from poking at a 2D screen.

It is software and technology associated with AI, including machine learning and natural language processing, computer vision, AR, Virtual Reality, and all other apps that support the creation and maintenance of a digital 3D world. We see great strides that will be made in Spatial Computing uses for many industry verticals, including **technology, media, and telecommunications (TMT)**, transportation, manufacturing, retail, finance, healthcare, and education.

Before we dig into everything that’s happening that caused us to write *The Infinite Retina*, let’s back up and think about the Prime Directive that is driving billions of dollars in human effort into Spatial Computing. Why do human beings need robots delivering food and building things in factories, and why do we need Spatial Computing devices on our faces so that we can work, entertain ourselves, educate ourselves, and collaborate with each other in new ways?

What is the Prime Directive? Does it have something to do with why humans spend more and more on technology or tools every year? Are any new trends, like our changing understanding of climate change, causing this change? Does culture itself change in a major way because of the Prime Directive?



Figure 1.1: Robert Scoble standing in front of a Tesla Cybertruck and wearing an Apple Vision Pro. Photo credit: Thomas Hawk.

What Makes Us Human?

Human beings are classified as *Homo sapiens*, which in Latin means “knowing man.” Modern *Homo sapiens* are believed to have appeared a little more than 300,000 years ago. The distinction between *Homo sapiens* and what came before has to do with the relatively sophisticated use of tools—tools that were used to survive more efficiently and with which humans gained control of their surroundings. Tools were also used by early humans to make art on cave walls and carve statuettes of female fertility goddesses. The tools served as augmenting devices—augmenting humans’ chances of survival and also of expression.

With modern humans, this augmentation can take the form of education, which in turn is used to gain knowledge. With knowledge, our chances of survival should be better. In many ways, our Prime Directive is to know how to better survive and how to better express ourselves by willfully creating and using tools for those purposes. It is a dual directive, for it cannot be proven that one gives rise to the other, but rather both are mutually beneficial. It is for both very practical and expressive reasons that tools have continued to be created from the time of early man to today. An example of a human being’s ingenuity that traverses both the practical and the expressive are the iterative inventions of the writing “pen and paper” combination. This combination tool, which goes back millennia, started out with cave walls, some form of patchworked dried grasses, as well as stone, serving as the “paper” and natural dye and sturdy reed, as well as a stone or metal chisel, serving as the “pen.” “Pen and paper” has been used to record both business and legal matters, as well as nonfictional and fictional narrative, and poetry, as well as visual art, such as paintings, when the “pen” is conceived as pigments. With the advent of the typewriter (and before that the Gutenberg press), there was even more of a separation between the practical and textually expressive and the visually expressive. A machine, the typewriter, was then replaced by the word processor and then the computer. These days, we utilize our computers and their smaller counterparts—the smartphone. Computers did not only replace typewriters; they are also in the process of causing people to question the continued existence of physical books and newspapers, as well as movie theaters.

Our Prime Directive to know how to better survive and how to better express ourselves now has a new channel—Spatial Computing. With Spatial Computing, the uses of the technologies of VR, AR, and AI eclipse those of the computers we know today. In the near future, we will no longer have to use a physical computer to do our work and browse the internet. We will also be able to do so much more with the three-dimensionality of Spatial Computing and speech recognition software. It turns out that our need to better express ourselves appears to include a need to experience a replicated reality.

Replicating reality in the forms of paintings, fiction, and films, as well as other forms, has existed as long as human beings have had the need to express the conditions of both their individual and social existence in an effort to better understand themselves. Experiencing a replicated reality also turns out to be a very good way to achieve a new skill and to get knowledge in general, and with the advent of Generative AI, replicating reality could now be a fast and almost effortless action. Spatial Computing is the next generation of imaging that is able to replicate reality, allowing the movement from 2D imaging to 3D. With 3D imaging, the replication of reality is able to be more closely related to the reality it is trying to represent.

Human beings seem to get satisfaction out of presenting and experiencing narratives that have the appearance of being real. An example of this is a movie. It is difficult to say exactly why we get such pleasure out of viewing a “good” movie. Perhaps it is empathy, but the question still remains: Why does empathizing with movie characters that appear to be real make us feel good, much less entertained? With Spatial Computing, the visuals are even more true-to-life and we are able to move through them (Virtual Reality) or incorporate and manipulate non-real objects into our real world (AR). AI adds another layer to the existing reality by organizing previously unconnected data into meaningful systems that could then be utilized in Spatial Computing to feed our Prime Directive needs.



Figure 1.2: Here, you can see the slums and other residential buildings as seen from the Four Seasons luxury hotel in Mumbai, India. Billions of people live in similar accommodations around the world and they will experience the world far away soon in Spatial Computing devices. Photo credit: Robert Scoble.

Drivers and Benefits

The benefits of Spatial Computing play right into our Prime Directive of knowing how to better survive and how to better express ourselves through our creation and use of tools. Our need to have replicated 3D worlds and objects in order to master our understanding and manner of expression is one that could be served by the software and technologies of AR, Virtual Reality, and AI.

Noted investor and Netscape founder Marc Andreessen has told markets he has a contrarian view of Silicon Valley's belief that AR represents a better investment than Virtual Reality. He noted that it is a privileged view—that in Silicon Valley, residents have tons of beautiful places within an hour's drive, from beaches to vineyards. Most people in the world, he said, don't have those advantages.

Walk through neighborhoods, even middle-class ones, and you will see millions living in small homes in high rises. Telling them that they will want to wear computing devices while walking through a beautiful area won't be hard. Instead, Andreessen sees a world where people will wear headsets to visit the natural beauty well out of reach somewhere else in the world. Even in the United States, only about 20 percent own a passport, so asking them to visit historic sites in, say, Egypt or Israel, won't be possible for most. We can, instead, take them there with Spatial Computing.

However, unlike Andreessen, we don't view the question of whether Virtual Reality or AR will be the "winner" between the two. We had predicted in our first edition that by the mid-2020s, our Spatial Computing devices would let you float between putting virtual things on top of, or replacing, things in the real world. With the Apple Vision Pro, and the Meta Quest 3.0, we solidly have that now.

Where we are going with this argument is that the hunger for these devices, services, technologies, and experiences that Spatial Computing affords will probably be far greater among the billions who can't afford a private jet to fly them to a Davos, Switzerland, ski vacation, or even afford a Tesla for a weekend jaunt to Napa or Yosemite. That seems to be Andreessen's greater point; that the investment opportunity here is grand because it not only improves the lives of billions but may lead to us saving ourselves with new education and new approaches to living, including being able to take courses using Virtual Reality and "travel" to different locations in the world without having to jump on an airplane.

However, an argument could be made that human beings are social—our social natures have aided us greatly in our need to survive and thrive—and that Spatial Computing is too isolating. With Spatial Computing, though, we could choose to experience and learn something solo or with others.

In our first edition, we said that Spatial Computing on its own can serve as the medium to interface with ideas, locations, processes, people, and AI characters, or as Betaworks' John Borthwick likes to call them, synthetic characters. These synthetic characters will replace the machine learning text bots that are currently ubiquitous—in effect, putting a 3D body to the words. The benefit of this is that we will feel like we are engaging with a real being that feeds our social natures. These synthetic characters are now called virtual beings and many of them are being powered by Generative AI.

Along these same lines, entertainment that utilizes Spatial Computing will make it even more true to life. Characters in these new kinds of narratives will be more real to us, allowing us to gain even more insights into the human condition. Spatial Computing is a major innovation; the latest in a long line of ideas and inventions aimed at improving the human condition. It makes our lives better, bringing knowledge to us faster and with less expense overall. Generative AI here, too, is having a big impact. Apps like Google Sora and those made by Runway and Pika AI are emerging at a very fast rate, to the extent that the entertainment industry will soon be disrupted by it. Soon, Spatial Computing will have much Generative AI content that is of good quality.

Potential Dangers

It could be said that with great knowledge comes great responsibility. There are several associated potential dangers that come with our use of Spatial Computing. Areas we will touch upon here include potential loss of personal control, dilution of the natural, and population segmentation.

In terms of potential loss of personal control, the major one that everyone has at the front of their minds because of Facebook's transgressions is the disappearance of privacy. Especially with AR, unauthorized use of data and media by companies or authoritarian use by governments could potentially be a problem. Location and spending data, along with videos made with the knowledge of the viewer and recording speech when headsets/glasses are worn, could present another wave of Google Glass-like uproar. However, we do not think this is going to happen since the uses of AR have been lauded over the last few years and there also seems to be more acceptance of glasses-like headsets due to the public's understanding that privacy issues will be addressed by AR hardware and software companies, along with much deeper utilities that today's technology affords. Companies will need to be especially clear as to what their data policies are and have appropriate opt-in policies to meet the expectations of the public. There is so much reward that could be received by providing data for AR purposes—these rewards should be heralded while understanding that there are those who would prefer to not share their data.

Another potential loss of personal control is one of over-advertising. As Keiichi Matsuda's 2016 six-minute nightmare concept film *Hyper-Reality* portrayed, a world where AR advertising is constantly overlaid over the real world is one that is unbearable.



Figure 1.3: Here is a still from Keiichi Matsuda's *Hyper-Reality*.

Having an opt-in system should solve this, but it might still be an issue when there is a willful exchange of some kind of visual goods with advertising. Companies will probably push the advertising threshold with potential consumers to see how far they can go in this case.

Talking about extremes, this brings up the point of addiction that has come so naturally with the advent of the smartphone. As with any other kind of addiction, it is certainly the case that personal loss comes attached to it. There are many cases of people who have died while taking selfies while in the throes of being distracted by their digital addictions. There is the possibility that a person walking around in a future blockbuster AR game will walk into traffic or even off a cliff, so technical safeguards based in computer vision will definitely have to be built in as an alert device.

Along with this kind of addiction comes a dilution of the significance of the natural objects and environments of the world as they appear in reality. This is relevant with regard to both augmented and Virtual Reality. Dystopian visions on this abound, ranging from people never wanting to be in the real world again, to the death of learning, to the abandonment of care about pollution and global warming. In Virtual Reality, a person could potentially hurt and even kill a digital character without having the full effect of what these actions would be like in real life.

A worry might be that these actions via augmented and Virtual Reality could become so common that the line between the imaginary and the real could become blurred to the extent that real people would then get hurt and killed. Industry oversight organizations might spring up to create ratings of Spatial Computing experiences that would rate the level of violent content so that viewing these experiences could be better managed and controlled. In this way, possible negative societal effects could be mitigated.

On the other side of this, the benefits of Spatial Computing will not be able to be shared by everyone on the planet due to economic reasons. Even in relatively well-off countries, there will be segments of the population that will not be able to afford Spatial Computing headsets or glasses. The result of this is that there will be great inequality with regard to information and productivity between the two groups. We believe that both laptop and desk computers, as well as tablets and phones, will eventually be replaced by Spatial Computing headsets or glasses. Already there's an entire new industry trying to make new AI-powered glasses, pins, and other devices that you'll wear around your neck or in your ears that bring various forms of Spatial Computing to us. Without Spatial Computing devices, both work and entertainment could prove to be difficult. Over time, we believe that the cost of Spatial Computing devices will come down due to technical efficiencies and product commoditization so that much more of the world's population will be able to afford them.

There are many uses for Spatial Computing devices, which is a main theme of this book. We will now provide a backdrop for understanding why Spatial Computing will have the impact on the world we believe it will.

Understanding the Natural World

Sometimes, the storm of change comes as a real storm. Visit Half Moon Bay, California, which is a sleepy coast-side town near San Francisco and Silicon Valley, and you'll probably meet "Farmer John." That's John Muller, and he moved there in 1947 to open a small family-run farm that is still operating, mostly growing pumpkins. We met him a while back and he told us that when he started his farm, he had to guess the weather, mostly by "feel." Today, he told us his son relies on satellite imagery and AI models to know what tomorrow's weather will hold, rather than holding his finger to the wind and trying to guess if a storm is on the way. By knowing the weather, he can better hire people to plant or harvest, and that knowledge saves his family a lot of money. Knowing a storm is coming a few days in advance can save much. Knowing that a heavy rainstorm that might flood fields is coming can save millions in seed costs, particularly in places like Kansas. Knowing a tornado will soon hit a neighborhood can save many lives.

Soon, Farmer John's family will be warned of changing weather with AR glasses aided by AI and computer vision, and they will be able to see storms in 3D instead of watching X (formerly Twitter) accounts like many farmers around the world do today (we built an X feed composed of meteorologists and scientists to watch our changing climate better at <https://twitter.com/Scobleizer/lists/climate-and-weather1>). The government and others have spent billions on this infrastructure. It is hardly the only expenditure humans have made to understand our changing environment.

In 2003, we had a discussion with Bill Hill, who was a computer scientist at Microsoft. He invented the font-smoothing technology that we use on all of our devices now, and he told us how he invented the technique. He had an interest in animals, thanks to reading tons of books when he grew up in poverty in Scotland and learned to track them through forests and meadows by looking at their footprints, or other signals. While doing that, he realized that humans evolved over millions of years to do exactly what he was doing as a hobby. Those who survived had deep visual skills and generally could see tons of patterns, especially in the green grass or trees that our ancestors lived in. If you didn't, you were attacked and eaten by an animal who was camouflaged in those trees. So, those who couldn't see patterns, especially in the green foliage around us, had their DNA taken out of the gene pool. He used that knowledge to figure out that he could hide visual information in red color fringes that surround each letter on your screen. Get out a magnifying glass and look at fonts on Microsoft Windows and you'll see his work.

Today, we no longer need to worry about a lion hiding in the grass waiting to eat us, but we have new challenges that continue to push us toward developing new tools that can help us continue this great human experiment that we are all in.

The lions of today might be those in the capital markets that are pushing our companies to be more profitable. Or, they might be a disability that is slowing down our ability to be our best. We think that's all important, but what did our ancestors do after they survived the day walking in the jungle trying to find food to feed their families? They gathered around fires to tell stories to each other after a day of hunting and gathering. We can see some of those stories even today as drawings on cave walls, where drawings show brave hunters facing down animals and dealing with the challenges of crossing mountains and streams. Today, we don't need to hunt to eat, but there are those who still want to do these things as sport, and what they are doing with technology demonstrates we are under a Prime Directive to build better tools and experiences, even to do something humans have done for many generations.

Let's look at another example of how people are utilizing technology in order to better understand their environment and augment their experiences. Vail Colorado resident, Ryan Thousand, is a senior advisor in emerging tech and healthcare innovation at KiloBiteData. During the day and many evenings after work or on weekends, he is a passionate fly fisherman who loves catching fish, not to eat, but to capture for his Instagram channel to brag to his friends before he releases them back into the stream for others to experience the same joy. He saw a need to track what kind of lures he was using as stream conditions changed, so now he's building 3D models of the streams he fishes in. While wading in the water, he wears smart boots that measure water flow, temperature, and other things. The data he captures using those and other tools is then captured on his phone and streamed to a database he's building, and also to the government, which is using the data from sportspeople to track environmental ecosystems.

His early attempts to build useful tools for fly fishermen brought him to Oakley's ski goggles, which had a tiny little monitor built into them. He sees a world, coming soon, where he can move away from having to view data through a tiny little monitor that gets in the way of the real beauty he is usually surrounded by, and instead transition to wearing a headset like Microsoft's HoloLens while fishing so that he can capture stream conditions on top of the actual stream. Soon, he says, people will be wearing a set of Spatial Computing glasses, which will include a virtual guide that will not just show you potential fishing spots in rivers and lakes before you go on a fishing vacation but will then help you learn how to properly cast a fly, and even tell you where to stand in the stream to have the best chance of catching a fish. Thousand's Prime Directive is to make fishing better and to help everyone who visits Vail learn more about the environment, which will help them save it for many future generations.

The changes, he told us, could change all human activity, even one as familiar as grabbing a fishing rod for a day of relaxation in a stream. That's due, in part, to things like new 3D sensors and new optics that will let us augment the real world with new kinds of digital displays.

Faster Wireless Leads to New Affordances

In talking with Anshel Sag, principal analyst at Moor Insights and Strategy (<https://www.linkedin.com/in/anshel-sag-7484127/>), it became clear just how far away 6G is, to come about sometime in 2029 or 2030. But we have 5G now, which is only halfway through its standardization and adoption. The 5G we have now will be much better toward the end of 2024 for Spatial Computing. The reason is a technical one; more operators are expected to adopt what's called "standalone 5G" or "advanced 5G," which is a 5G frontend with a 5G backend, rather than "non-standalone 5G," which is a 5G frontend with a 4G backend. This is included with release 18 for 5G.

Specific enhancements with regard to release 18 of 5G and Spatial Computing include power savings and capacity improvements, along with the optimizing of experience for high-bandwidth **extended reality (XR)** applications, and an enabling of what is called “REDCap” at low power. REDCap is basically 5G but cut down to a minimum capacity; REDCap requires standalone 5G.

With the power savings and capacity improvements that release 18 brings, Spatial Computing devices will be able to render experiences more efficiently and with lower latency. This will be a great improvement and what Anshel Sag says “is going to be a big story” and what everyone will be talking about.

Not only will release 18 be especially good for Spatial Computing, but it will also be good for AI, which Spatial Computing utilizes. Most AI is currently held and calculated in the cloud, which is very expensive. In order to enable hybrid AI, which allows for AI to be held and calculated locally, as well as in the cloud, using edge computing, you need a fast network that moves these models. Sag believes that advanced 5G will be a huge enabler of AI and Spatial Computing. He says, “I think this is maybe an opportunity this year and next to really actually see more mature implementations of all three together.”

In Spatial Computing, AI is used for computer vision, personalization and virtual assistants, and the creation of visuals via Generative AI, along with a slew of other uses, so improvements regarding the local implementation of AI will have very positive effects.

On the home and office front, Sag says we should expect Wi-Fi 7 this year, which enables 2.4, 5, and 6 gigahertz to be used together at the same time, as well as allocating power efficiently across the three. In the future, he sees more XR headsets using smartphones for their compute power and believes Wi-Fi 7 will be a big improvement for Spatial Computing. This, along with release 18, then 19, and 20 for 5G, will have a huge impact on Spatial Computing and will change how we run our businesses, how we shop for new things, and the games we play with each other.

Data, Data, Data Everywhere

Data clouds and data floods are coming to our factories and businesses, thanks in part to the 5G infrastructure that’s been built. This is bringing a storm of change and a need to see patterns in that data in a whole new way.

A few years ago, we visited the new Jameson Distillery near Cork, Ireland. At one point, we asked the chief engineer, who was proudly showing us around its new building and machines that make the whiskey, “How many sensors are in this factory?”

“So many I don’t even know the number.” That was more than nine years ago. Today, some warehouses and factories have hundreds of thousands of robots and millions of sensors. You won’t find anything useful in the data that’s streaming off by using Microsoft Excel; there’s just too much data to look through grids of numbers.

Los Angeles-based Suzie Borders has a better idea. Turn those millions of numbers streaming off of sensors into something that humans can better make sense of; for example, a simple virtual light on top of a factory machine. She showed us what she meant by putting us inside various devices, including a Magic Leap Spatial Computing headset and a variety of Virtual Reality headsets, where she walked us around datasets. Datasets that don’t look like datasets at all! They are more like a new kind of virtual interface to data that you can grab and manipulate.



Figure 1.4: Suzie Borders, CEO/Founder of BadVR, wearing a Meta Quest 3.0. Photo Credit: BadVR, Inc.

Over the last few years, Borders has seen an increased focus on data visualization for specific industries and use cases, with more emphasis on solving problems versus just visualization, along with the use of AI to collect and analyze specific data. A benefit has been to visualize more data on headsets without being tethered. This is all in line with what we are seeing in terms of Spatial Computing being combined with the use of AI. All of this bodes extremely well for when AR glasses become the norm.

Yes, people will resist wearing glasses, especially if they don’t wear prescription glasses, but those who dive in will find they get the raises and kudos for seeing new ways to make companies more effective and efficient, and we are betting a lot of those new ways will be seen using Suzie’s software.

The upgrade from 2D technologies to Spatial Computing's 3D technologies has its roots in the analog to digital transformation that happened prior. It is important to understand some of the impetus behind the prior change and what significance it has for our shift.

Electric Sheep Robotics Dreams

Paradigm shifts, which come in storms of new technologies, maim big companies and create new stars. The iPhone marginalized both Eastman Kodak and Nokia. Both companies were dominant at one point; Kodak even invented the digital camera, but new technology created new stars and took away the relevance these once had.

The same will happen in the Spatial Computing decade, which is just beginning. The storm will be violent and swift but to understand why this is so, we need to go back to that earlier time.

In this vein, Nag Murty and his team at Electric Sheep Robotics are disrupting the worlds of landscaping and lawn care.

His company has built two robots, so far, with more on the way, that take care of various landscaping jobs. It bought several landscaping companies to get the data, and understand the customers, before automating them with robots.

His robots, trained on that data, use computer vision in a unique way. As we did the research for this book, we visited 11 states in 2023 and while driving by thousands of homes, we saw many lawnmowing robots out front. These existing robots are fairly dumb, though; the homeowner needs to put a wire around their lawn so these robots know where the edge of the lawn is, which turns them around to cut another row of grass.

Electric Sheep's robots are different. They need no wire, which lets them be placed on a lawn they never have seen before and get to work, just like a human. Each has a camera and makes a new path, just like a Tesla autonomous vehicle can drive down a street it never has seen before. The AI/computer vision can "see" the lawn via a camera on the front of the device and plot a path, just like an autonomous vehicle does.

When we visited, they were putting the finishing touches on their new robot, which can edge the grass and do other tasks with a replaceable arm. The staff were training it with NVIDIA simulation software. Laying down a fake tree in the simulator, or a bush, or a bump, they could train the robot on what to do with these new obstacles.

Then, they could take it to a fake lawn they have set up in their San Jose, California, headquarters, before taking it out to a real lawn (their headquarters has huge lawns surrounding their warehouse, which offers lots of testing grounds; plus, their first clients include some big hotels and golf course businesses).

It also has language recognition capabilities as a UX to be used to heed voice commands. Sheep Robotics is a Series A company, with four locations across the US. “Our approach is that the fundamental problem to solve is data and robotics,” says Nag.



Figure 1.5: Two of Electric Sheep Robotics' autonomous lawnmowers doing their jobs. Photo credit: Electric Sheep Robotics.

Electric Sheep rents out its lawnmowing robots to commercial customers. This seemingly mundane company is following the route that Elon Musk says he is taking with robotaxis as a service. We have long advocated that this would be what Elon Musk plans to do with his Teslas. Here, Electric Sheep is following his stride. Commercial landscaping is a \$600 billion business. Electric Sheep offers landscaping services to commercial customers and offers their Electric Sheep robots to rent out as a part of their services.

Their system does not require boundary wires as the “robo-sheep” can sense when there is an edge and can make out whether there is a shrub, a sprinkler head, or humans in front of them. There is no pre-made map—the device maps out its environment as it goes along.

Nag quips that there is “one model, many different bodies,” in that the electronics that run the autonomous lawnmowers could be swapped out to directly work on another kind of robot. This allows their company to scale in a very smart way.

Electric Sheep is emblematic of companies that have hardware that uses computer vision, which is a key element of Spatial Computing. AI is used by computer vision, as well as the language recognition that Sheep Robotics offers. So, the elements of AI and computer vision are very important to Spatial Computing.

Dancing into Different Worlds

The storm of change will mean that the nerds who used to build the computer software of old, by typing thousands of lines of code into a black box on a flat screen, might need to learn to dance.

Yeah, if you watch someone play in Virtual Reality, you might think they are dancing, but that isn’t what we are thinking about. There’s a new type of computer science underway: one that uses choreography to train the AI systems that control autonomous cars, robots, and virtual beings, and even present humans in augmented and Virtual Reality with interfaces that better serve us.

Machines can be cold. In the worst of cases, they can crush us on factory floors, or kill us in the streets, like a computer-controlled car developed by Uber did one unfortunate night in Arizona. In the best of cases, they can give us superpowers and make our lives easier, but even then they often can be made to better serve humans in a dance, of kinds.

Did you know there is an entire conference for research on these new ways of training machines called “Choreographic Interfaces?” We didn’t either until we were introduced to Dr. Catie Cuan, who received her PhD at Stanford University in 2023 on the topic. As she showed us around Stanford University’s AI and robotics labs, where she is working on her studies, we discussed how innovation teams will soon be changed by her, and other people who are using techniques that humans use in dance choreography to train computers.



Figure 1.6: Dr. Catie Cuan tells us about using dance choreography to train robots in the lobby of the Bill Gates building at Stanford University, which is the building where Google started.

Photo credit: Robert Scoble.

What makes her unique among most of the computer scientists we've met? She's a former dancer. "Dance and choreography is all about moving with intention," she says, pointing out that she spent more than two decades dancing and that she uses the knowledge she gained by studying how human bodies could be made to move in different ways, and then translates that training and knowledge to teaching robots and virtual beings to move.

Cuan worked for automobile companies that are developing autonomous cars, among others, as part of her studies (she did some of her schooling at the University of Illinois, which is where Tesla and PayPal started, among others). She told us that humans go through an intricate dance of sorts as they wave each other through intersections, for instance, and cars without human drivers can't communicate with humans on the street or in other cars that way; so, they need to be trained both how to recognize human gestures, say, like a police officer standing in the middle of an intersection directing traffic, and communicate with humans their intention to move or stop, but do so in a pleasant, human way.

Now, as computing will be everywhere, we need new interfaces. Ones that respect us, understand us, warn us properly, all in a very human way. Building these new interfaces requires new kinds of innovation teams with, even, a dancer or two to help out. Why? Well, she says that she sees an optimistic future for human/machine interactions, and what better way to make the machines more fun and engaging to be around? As this storm of change comes, though, it may frustrate many people who want things to stay the way they are, or who can't wrap their heads around the fact that they need to work with someone doing choreographic work. Many pioneers face similar resistance in their careers. We suggest that if you feel yourself resisting this change, you might need to be the one to change lest you, or your company, be left behind.

The Frustrated Pioneer

Doug Engelbart had a frustrated spirit when we first talked with him in 2005. If you search YouTube for “mother of all demos,” you’ll find the video of him demoing something that looked like the Macintosh. He guided a mouse cursor across the screen and showed off many new computing concepts that would dramatically change how we view and use computing. The official title was “A research center for augmenting human intellect,” way back in 1968. Almost 20 years later, the Macintosh was born from many of the ideas he demonstrated in that demo. He later went on to win many awards and accolades, but he was particularly proud of the National Medal of Technology, the United States’ highest technology award, yet here, he was sharing his frustration that he wouldn’t be around to help humans see what he told us would be his most important work: augmenting human beings. In other words, giving them the tools to make themselves better.

“Why are you frustrated?” we asked him. That led to a discussion of how Engelbart was kicked out of his own research lab back in the 1970s (which is the same lab, SRI International, then named Stanford Research Institute). This is the lab that brought us Siri and HDTV, and played key roles in many, many other things, including the internet itself). Some of his coworkers say that he was too focused on the future and not on getting revenue for the lab, but Engelbart answered, “People don’t understand what I’m saying.” He explained that he was dreaming of a better future, where humans and machines would join together.

In one trip, he pulled out a glove that would let people touch and type faster than on keyboards.



Figure 1.7: Robert Scoble's son, Patrick, talks with Douglas Engelbart and plays with the original mouse. Photo credit: Robert Scoble.

When asked why they didn't understand, he answered that for that, you have to go back to the context of the day. Back in the 1970s, before personal computers came along, computers were run by either data entry clerks or by PhDs and other highly trained people who wrote the code that ran them. When he told his coworkers and others that someday we'd have a supercomputer in our hands, he was laughed off as a crackpot.

He had always dreamed of a world decades away and talked with us about how computers would soon lead to different realities. As the inventor of the original mouse, he was always looking for ways for humans to use computers better than typing on a keyboard or stacking piles of punch cards into trays, which is what most computer users back in the 1960s were doing. Today, we all use the progeny of his work, and this pioneering spirit of dreaming of a better world, where humans would have better control of the technology they invented, continues to drive us to new and greater innovations.

His idea was that computing could go beyond merely performing calculations and that it could be used to augment the capabilities of the human mind. Why was he frustrated? Because even as he was in the final phases of his life, he could see a world where people would see or sense computing all around them and be able to touch it, maybe with gloves, but hopefully, he told us, with sensors that could see the human hands, eyes, and maybe someday into their minds.

These are ideas that even today aren't largely used or popular. Virtual Reality has only sold a few million headsets at the time of writing this book. Most humans haven't had even an hour in one of those, and here this man is telling us about a world he saw way past Virtual Reality to one where computing was on every surface and where you could use your hands, eyes, and voice to control it! His thinking on this topic led him to develop what came to be known as Engelbart's law: that computing was increasing at an exponential rate, so we would be able to exponentially increase our performance as well. We are seeing this become one of the key drivers behind the development of the technology we see about to disrupt seven industries, and probably more.

He saw a world where we would soon have technical capabilities beyond the imaginations of most people; for instance, cars that could drive themselves, or systems that would let us see a new digital layer, making our understanding of the world far better than before. His ideas of augmenting human performance through technology changed all of our lives, and soon will bring us both new powers and a new understanding of what it means to be human, but to do that, he needed to wait for our machines to get new senses. We wonder if even Engelbart himself, though he usually was decades ahead of others in his thinking, could have imagined the 3D sensors that are now in our phones and cars. Could he imagine the databases and cultural changes that they are now bringing? Knowing how he thought, he might resist some things for a few microseconds but then he'd smile, lose his frustration, and see that someone else is pushing the world toward his ideas. Let's meet one of these people now.

Breakthroughs in Sensors

David Hall had a building to fill. In the 1980s, the company he started, Velodyne, built subwoofers. Back in the day, they were the best subwoofers because they were the first to be digitally controlled. Audio fans lined up to get one since they made your audio system dramatically better. However, by around 2000, he realized that China had taken over manufacturing, so he sent his production to China since he couldn't get parts anymore to build them within the United States, as he told us in an interview¹.



Figure 1.8: Velodyne founder/former CEO, David Hall, shows off his LIDARs and talks about his career as an inventor of everything from new kinds of subwoofers to self-leveling boats.

Photo credit: Rocky Barbanica.

That left him with an empty building, so he started looking for new things to do. He heard of the **Defense Advanced Research Projects Agency (DARPA)** Grand Challenge, a prize competition for American autonomous vehicles, funded by the DARPA, a prominent research organization of the US Department of Defense, and readied a vehicle to enter to attempt to win the \$1 million prize. The military wanted to spur the development of autonomous vehicles. Hall soon realized he didn't have the software engineers needed to win because he was against teams of university students from schools like Stanford and Carnegie Mellon, but that showed him a new market: making sensors for the other teams. He had the engineers and production ability, thanks to his empty building, to do that, and so he started building LIDARs.

LIDAR, which stands for light detection and ranging, is the spinning thing you see on top of autonomous cars often seen driving around San Francisco or Silicon Valley. Most of them are from Velodyne. Google's early cars used his LIDARs, which spun dozens of lasers many times a second. The cars used that data to build a 3D model of the world around the car. "If they have a spinning one, that's mine," he told us.

It is worth mentioning that LIDAR sensors could even be found in the iPad Pro and the Pro model of iPhone, as well as some Android phones, and these improve their AR capabilities.

Invisible light from his spinning device bounces off the road, signs, other vehicles, and the rest of the world, or in the case of his boat, waves in the ocean. A computer inside measures how fast the light returns to the sensor and makes a 3D model of the world. Without these sensors, the car or robot can't see, so it can't navigate around the world.

His new technology plays an important role in Spatial Computing because it gives computer scientists new data to train new kinds of systems. Google Maps, for instance, is better, its former head of R&D, Peter Norvig, told us, because, as it used Hall's LIDARs, it trained a system to gather data about street signs the car was passing.

"(Starting a LIDAR company) was over everyone's dead body," he said, because his company was so focused on making subwoofers. "I look for things that are electrical and mechanical." He convinced his team that they had a new market and a new thing to do, and now is an important player in the autonomous car space. He also used them to build a boat that won't make you sick (computers control hydraulic lifts to glide over waves).

In 2019, Google also used that data to turn on new AR navigation features, where data from your phone's camera was used to compare with 3D models of the street built using data gathered from one of Hall's LIDARs. It is this impulse to build machines that can better see, and hence, make humans more powerful, that enabled these new features and for that, Hall is an American engineering star who is bringing us a whole new world. As we're about to see, this sort of development was not only happening in America, with sensors attached to a Spatial Computing device much smaller than an autonomous car.

Thousands of miles away, in Tel Aviv, Israel, was another innovator working on similar 3D mapping technology, who had a similar philosophy about how 3D sensors would soon play a key role in making human life better. That man is Aviad Maizels, and he started PrimeSense. PrimeSense was purchased by Apple and did much of the engineering on the sensors inside modern iPhones that see your face in 3D, which lets them unlock just by seeing your face, and soon will do much more.

Maizels told us back in 2013 that 3D sensors would soon be on everything. We met Maizels at the company's first suite at the Consumer Electronics Show in Las Vegas, where his company showed what 3D sensors could do—everything from seeing you buying cereal boxes to drawing on a standard desk with your finger. Today's Amazon stores demonstrate his vision was correct. In them, sensors watch you take products off the shelves and properly charge you with no registers or lines to wait in, saving you time and hassle.

It is worthwhile, though, to go back to that booth and take a look at how large PrimeSense's sensor was back then. In just a few years, it would shrink 20 times down to fit in the notch on current iPhones. Soon, it may disappear altogether since standard cameras now can build 3D models of the world, which has been proven by companies like the recently acquired 6D.ai, which can do much more with a standard camera than even PrimeSense demonstrated in that suite back in 2013 in this video: <https://youtu.be/4VtXvj4X0CE>.

Even in autonomous cars, this argument between using standard cameras and 3D sensors is still playing out. Tesla's engineers told us that they are betting on cameras instead of the more expensive LIDARs that companies like Velodyne are producing.

If you visit Silicon Valley, you will see lots of autonomous cars from companies like Cruise, now owned by General Motors, Waymo, formerly of Google, and others, all using LIDAR to sense the world around the car, while Tesla is betting on cameras, along with a few cheaper sensors like radar and ultrasonic sensors.

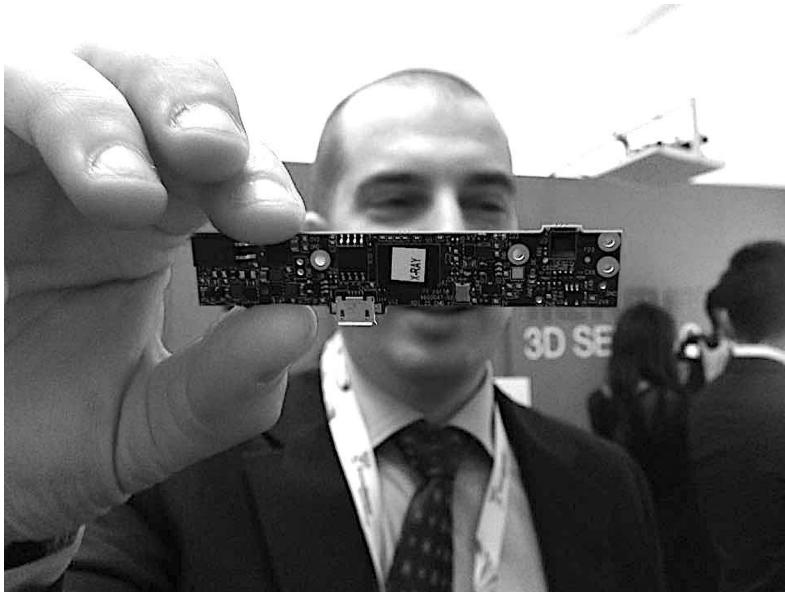


Figure 1.9: PrimeSense founder/CEO shows off the 3D sensor he and his team invented. This technology has now been sold to Apple, which shrunk it and included a similar sensor on the top of modern iPhones, which senses your face to unlock your phone or do AR “Animojis.”

Photo Credit: Robert Scoble.

That is a tangential argument to the one that we are making. These innovators brought our machines new ways to see and because of that, we now have new ways to see ourselves and the world we live in, which will grant us superpowers, including ways to be remembered by future generations that only science fiction dreamed about before.

Driven to Improve

“Watch this,” Elon Musk, CEO of Tesla, told us as he raced Jason Calacanis, driving a top-of-the-line Corvette through the streets of Santa Monica in the first street drag race done in a Tesla on February 18, 2008.

Good thing the video hasn't survived since speeds in excess of the speed limit were quickly reached and Calacanis was left in the dust (he soon became an early investor in Tesla, and owner of his own Tesla). Elon was behind the wheel of serial model #1 of the first Roadster Tesla produced. What we learned on that ride has become a deep part of Tesla's story and gave us insights into Elon's philosophies early on. Yes, he was demonstrating the advantages of electronic motors. He bragged that a Tesla, with its high-torque electric motors, could accelerate faster from 0-60 mph than his million-dollar McLaren F1.

He bragged about something else, though, that stuck with us all these years: electric motors could be programmed never to slip, unlike most gas engines, which have to "wind up" 400+ parts to apply torque to the pavement. We got a look at this recently as the co-author, Robert Scoble, drove his Tesla Model 3 through the snow in Yosemite. It never slipped, even while going uphill in icy conditions.



Figure 1.10: Elon Musk, left, and Jason Calacanis check out the first Tesla only days after Elon got it off the factory floor, before heading out for a street race where Elon demonstrated how much faster it was than Calacanis' new Corvette. Photo credit: Robert Scoble.

Today, Elon is bragging about Tesla's safety: every new Tesla comes with nine cameras, along with several other sensors, and he is adding new self-driving features every month to the fleet of hundreds of thousands on the road, thanks to an over-the-air update system that is way ahead of other automakers. In early 2019, he announced the production of a new kind of machine learning chip with 21 billion transistors on board; he says these will make Teslas the most advanced cars on the road and will make them much safer than those that only have a human to steer them.

It's called Autopilot because, as of 2019, a human is still needed to drive the car in city traffic and needs to be available in case other situations arise. It doesn't navigate around new potholes, for instance. That's a temporary step, though, as Elon has demonstrated full self-driving features where the car can drive without human intervention.

Industry observers like Brad Templeton, who worked on Google's autonomous vehicle program and who also owns a Tesla Model 3, say that Elon is too aggressive with his timeline and that it may take most of the 2020s to make it safe enough to remove humans from the equation completely, although there are currently some fully autonomous cars in action in some US and Chinese cities. As of 2024, Tesla's **Full Self-Driving (FSD)**, still isn't good enough to drive around without a human in the car but Elon Musk says he will soon show off its plans for a new "Cybercab" that will drive people around with no human driver.



Figure 1.11: Some of Tesla's Autopilot/Full Self-Driving Programming team hang out with Scoble's son at a Salinas Tesla Supercharger and talk about the future of autonomous cars.

Photo Credit: Robert Scoble.

What does this have to do with Spatial Computing and especially our Prime Directive? Everything. The techniques Tesla's engineers are developing (along with others like GM's Cruise, start-ups like Zoox, or the former Google team that now is starting a company called Waymo) are very similar to the techniques that Niantic uses to enable a virtual Harry Potter to walk around the real world and are similar to the techniques that Apple, Microsoft, Facebook, and others will use to guide users around the real world in Spatial Computing glasses that will soon come to the market.

Teslas are already saving their users' lives, as you can see on YouTube, as Teslas automatically stop or drive around potential accidents, even at high speeds. Don't discount the many hours given back to their owners as cars automatically drive in traffic during commutes. Those hours represent lives, too. Hours, er, lives, that can be used to do other things.

In *Chapter 4, Vision One—Transportation Automates*, we'll dig deeper into the other changes that soon will come as humans won't be needed for cars to move around as the machine learning/AI that runs the self-driving technology is put to work doing other things.

Things like going to an Amazon distribution center to pick up packages in a much more secure way than having a delivery person leave them on your front porch, where they are easy to steal. Even in 2019, the computers in a Tesla also watch for people tampering with the car and can automatically record, say, someone breaking a window to steal something inside and share that with owners via an app on mobile phones.

Hands-On Use

While Elon, and his teams, were working to make cars using machine learning to look at the road ahead, Andy Wilson was working in his lab inside Microsoft Research to come up with better technology tools, not for our roads, but for our hands.

Wilson was the first to show us how a computer could watch your hands for gestures. He demonstrated how a computer could see that you had pinched your thumb and finger together.

His first demo showed how you could "grab" a map in midair and zoom it and twist it. This gesture survives even more than a decade later on Microsoft's HoloLens.

Wilson has been trusted for decades by Microsoft. Back when Bill Gates was giving speeches, Wilson built demos to show off the most advanced technology. That earned him his own AR lab. Once when we visited, he had a set of projectors and sensors hung from tresses overhead. He demonstrated how he could put computing on every part of your body, and even the room. "Here you go," he said, while dragging a virtual photo out of his pocket and handing it to us.

It was the first time we had experienced Spatial Computing where we were “inside” the computer and it was interacting with us.

That research, along with others, came together to build the basis of HoloLens. We first started visiting Wilson’s lab back in 2005, but there’s a guy who has been working on this stuff for far longer: Tom Furness.

His work in the 1960s was as an airplane cockpit developer, and that’s when he built new virtual interfaces that we think of today as Virtual Reality or Spatial Computing. Now, he is a professor emeritus at the University of Washington and former founder of the Human Interface Technology Lab. That school did seminal work on studying how Virtual Reality could be used to cure pain. Its results showed that Virtual Reality, when used with burn victims, is better at removing pain than morphine is. This is something we hope becomes more adopted as we try to keep people from being addicted to opiates.



Figure 1.12: Tom Furness, right, and Robert Scoble hang out at Digital Raign’s Reality Summit.

Photo Credit: Robert Scoble.

He told us that he and his students had been developing Spatial Computing to engage, enlighten, and even entertain people in new ways. We include him here not just because he is seen as the grandfather of Virtual Reality, but because he is a great example of the pattern of how innovation frequently happens: first for military uses, later moving to consumer applications. Many of the industry’s leaders follow this pattern and that is why technology centers tend to cluster around strong military centers, whether in Tel Aviv, Silicon Valley, or in Las Vegas, where the military tests out new kinds of airplanes and continues Furness’ early work today.

We visited Nellis Airforce Base, where we met several pilots who fly F-35s. The pilots in them use a big, expensive, AR headset. One of them said something that stuck: “I’ll never lose to an F-16. Why? Because I can see you and you can’t see me, and I can stop.” The F-35 not only has those magic headsets, which let pilots see the airspace in great detail, but are designed to be stealthy against radar, so they can’t be detected by other planes. They are equipped with engines that can direct the flow from the jets to literally stop in mid-air, which the F-16s can’t do.

It isn’t lost on us that recently there’ve been reports that newer F-35s get rid of the pilots altogether and that the AI flying usually beats the pilots because the computers flying can perform maneuvers that humans just can’t handle due to the forces involved.

Wars handled by machines is a controversial thing to be certain, and one we won’t take sides on here, but that means fewer pilots coming home in a box due to accidents or being shot down, and the military sees that as just as good a thing as seeing fewer deaths on the road. No matter what way you look at these new technologies and their role, you can’t deny that most of these technologies were first designed with a military use in mind.

Beginning

In this chapter, we have thrown light on just a portion of the stories of innovators and the innovations that they are bringing to market, all of which we will all use by 2030, or most likely by 2025. We’ve given you an initial indication of how versatile and useful Spatial Computing is—for use by farmers, fly fishermen, businesspeople who use data analyses, to people who ride in Teslas and, soon, autonomous vehicles. Technologies such as machine learning and 5G have been further strengthening the reach of Spatial Computing.

We’ve been telling people we are just at the beginning of a new 50-year product cycle, which brings a new kind of personal computing that many have said will lead to the singularity—a deep merging of humans with machines. Already in R&D labs, researchers are showing a new kind of computing coming where you will just “think” of doing something, like sending a message to a loved one, and a message will appear in mid-air and you’ll send it just by thinking “send.”

The road to Spatial Computing has been a very long one—one that humans have been on ever since they started making tools that could substantially improve their chances of survival and their quality of life. Where we stand now is in an iterative stage that will yield significant gains. It is as if we are back in 1976 and Wozniak and Jobs just showed us the prototypes of the Apple II. We are still seeing improvements to personal computers 40 years later; however, this new type of computing will be a major leap.

It appears that the closer a created Spatial Computing experience is to real life, the more satisfaction a viewer has from it—both from practical and entertainment perspectives. The utility makes sense given the nature of Spatial Computing. In which way and how far humans will take Spatial Computing is very much the reason for this book. We are approaching it from an enterprise perspective, focusing on the particulars of industry vertical use, although our discussion of particular technologies will make it clear that there is a massive potential for wide consumer use of Spatial Computing.

For the rest of this book, we will detail and discuss what we think will happen in Spatial Computing within the next five years, give some major indications of what could happen in 10, talk about what it could mean for humanity in 25 years—focusing on culture and society—and give some far-reaching comments on what we might see in 50 years. We will look at seven industries that are about to undergo massive changes and radical disruptions and to understand why, we'll dive into the technology driving this perfect storm of change and take a look at some of the changes coming to each of them.

References

1. David Hall's interview with Robert Scoble: <https://youtu.be/2SPnovjRSVM>

2

Four Paradigms and Six Technologies

You aren't reading this book to learn about small new features of existing apps, but rather to more deeply understand what we are calling the "Fourth Paradigm" of personal computing—Spatial Computing.

This chapter will bring you up to date on the technology you need to know about to understand where Spatial Computing came from, how it's integrally related to AI, where it's going, and why it's so fundamentally different from the desktops or mobile phones that came before, even though it uses a lot of technology invented for the machines of earlier paradigms. Spatial Computing is the fourth paradigm, and we see it as the most personal of all the personal computing paradigms that preceded it. We will cover six technologies that enable Spatial Computing to work: Optics and Displays, Wireless and Communications, Control Mechanisms (Voice and Hands), Sensors and Mapping, Compute Architectures (new kinds of Cloud Computing, for instance), and AI (Decision Systems).

The Four Paradigms

What do we mean by "paradigm"? The evolution of the user experience is the easiest way to understand this. The first paradigm has to do with the beginning of personal computing and the interface being text-based. With the second paradigm, graphics and color capabilities were included.

Movement and mobility were added with the third paradigm in the form of mobile phones. With Spatial Computing, the fourth paradigm, computing escapes from small screens and can be all around us. We define Spatial Computing as computing that humans, virtual beings, or robots move through. It includes what others define as ambient computing, ubiquitous computing, or Mixed Reality. We see this age as also being inclusive of things like autonomous vehicles, delivery robots, and the component technologies that support this new kind of computing, whether those are laser sensors, Augmented Reality (including machine learning and computer vision), or new kinds of displays that let people view and manipulate 3D-centric computing.

We will see seven industries that will be dramatically upended by the 3D thinking that Spatial Computing brings and go into more detail on the impacts on those industries in further chapters. Soon, all businesses will need to think in 3D. This shift will be difficult for many. Why can we say that with confidence? Because the shift has started already.

The Oculus Quest, which shipped in May 2019, was an important product for one reason—it proved that a small computer could run an on-the-face spatial computer that brings fantastic visual worlds you can move around in and interact with. It proved that the world of Spatial Computing had arrived with a device many could afford and that is easy to use. It was based on a Qualcomm chipset that is already years old as you are reading this—all for \$400, and with controllers that could move freely in space and control it all! Now that our bodies can move within a digital landscape, the Spatial Computing age, that is, “the Fourth Paradigm of Personal Computing,” has arrived and will only expand as new devices continue to arrive, such as the Apple Vision Pro. What we are seeing now are the Spatial Computing counterparts of products like the Apple II or the IBM PC of earlier paradigms.

The Quest isn’t alone as evidence that a total shift is underway. Mobile phones now have 3D sensors on both sides of the phone, and brands are spending big bucks bringing Augmented Reality experiences to hundreds of millions of phones and seeing higher engagement and sales as a result.

Look at mobile phones and the popularity of Augmented Reality games like *Minecraft Earth* and Niantic’s *Pokémon GO*. You play these while walking around the real world. A ton of new technologies are working in the background to make these 3D games possible. And with the advent of the Apple Vision Pro and its passthrough AR technology, these games could be played on that device.

How is this fourth paradigm of computing different? It is additive, including the technologies that came before it, even as it brings quantitatively different kinds of technology to our eyes and world. Spatial Computing promises to do what the previous three paradigms have failed to do: to map computing to humans in a deep way. Elon Musk says it will increase our bandwidth; what he means is that our computers will be able to communicate with each of us, and us with our technology, in a far more efficient way than ever before. Paradigm one kicked it off by enabling us to converse with our own computer, one that was in our homes for the first time, through a keyboard. That brought a revolution, and the next paradigm promises a huge amount of change. However, to understand the amplitude of that change, we should look back to the world that Wozniak and Jobs brought us way back in 1977.

Paradigm One—The Arrival of the Personal Computer

The Apple II is as important as the Oculus Quest, even though most people alive have never used one. The Quest brought a new kind of Spatial Computing device to the market—one that mostly was biased toward Virtual Reality, where you can only see a virtual world while the real world is hidden from view. This \$400 device was the first from Facebook's Oculus division that didn't require sensors to be placed around you or a wire from the headset to a PC. It was all self-contained and it powered on instantly, which dramatically increased usage numbers.

Where the Quest let everyday people think about owning a VR headset for the first time, the Apple II acted in the same way for those back in the late 1970s, allowing people to own a personal computer for the first time! As the 1980s began, a lot of people, not just governments or big businesses, had access to computers. Four decades later, we are seeing the same trend with Spatial Computing.

The Apple II, and later the IBM PC, which ran Microsoft's DOS (Disk Operating System), introduced the integrated circuit into a consumer-focused product and enabled an entire industry to be built around it. By the end of the 1980s, literally every desk inside some corporations had a personal computer sitting on it, and the names of those who made them, such as Michael Dell, are still well-known today.

We expect that Spatial Computing will similarly lift previously unknown people up to wealth and household name status.



Figure 2.1: Steve Wozniak stands next to an Apple II at the Computer History Museum. He is the last person to have the design of an entire computer, from the processor to the memory to the display drivers, in his head. After this, teams took on the engineering of all those parts.

Photo credit: Robert Scoble.

Speaking of Dell, he may not be the one who gets credit for starting the personal computing age, but that's only because he was able to make personal computing more accessible to the masses. He did this by licensing Microsoft's operating system and putting it in lower-cost machines, allowing companies like Dell to make many sales. We foresee that a similar pattern will probably play out in Spatial Computing. Apple entered the Spatial Computing market with its Apple Vision Pro in February 2024, which is much more expensive than, say, competitive products from Google and Meta, not to mention those from China.

It is hard to see how early computers that could only display black-and-white computations on a screen could be relevant to Spatial Computing, but the concepts that those computers introduced remain relevant even today, such as files, printing, saving, and others. The thing is that back in the beginning, there were no graphics; computers were way too big to hold (much less dream about putting into your pocket) and much harder to use. Those of us who learned computing way back then remember having to look up commands in manuals to do something simple like print, and then having to type those commands onto the screen.

Spatial Computing, or computing you can move through, is actually joined by much-improved voice technology. Companies like Otter.ai are understanding our voices, and systems that are now being integrated with AI like Apple's Siri, Amazon's Alexa, Google's Assistant, and others are waiting for us to speak to them.

Currently, you can just say something like "Hey Siri, can you print out my report?" and it will be so. This new world of convenience that is being ushered in is, in our opinion, far preferable to the days of code and command lines that we saw during the first days of personal computing!

The first computers arrived with little excitement or fanfare. The first ones only had a few applications to choose from, a basic recipe database, and a couple of games. Plus, Apple founders Steve Wozniak and Steve Jobs were barely out of high school, and the first machines were popular mostly with engineers and technical people who had dreamed of owning their own personal computers. Those days remind us a lot of the current Virtual Reality market. At the time of writing this book, some millions of VR machines have sold. In its first year, only a few thousand Apple IIs had sold. It was held back because the machines were fairly expensive (in today's dollars, they cost more than \$10,000) and because they were hard to use; the first people using them had to memorize lots of text commands to type in.

It's funny to hear complaints of "There's not enough to do on an Oculus Quest 3," which we hear frequently, as well as the relative dirth of content for the Apple Vision Pro. Hello, you can play basketball with your friends in Rec Room! Try going back to 1977, when the first ones basically didn't do anything and, worse, you had to load the handful of apps that were out back then from tape, a process that took minutes and frequently didn't work right at all. Wozniak told us his wife back then lost her entire thesis project on an Apple II, and even he couldn't figure out how to save it. Those problems aren't nearly as frequent in the days of automatic saving on cloud computing servers.

Why is the Apple II, along with its competitors and precursors, so important? What was really significant was that people like Dan Bricklin and Bob Frankston bought one. They were students at Harvard University and saw the potential in the machine to do new things. In their case, they saw how it could help businesspeople. You might already know the rest of the story; they invented the digital spreadsheet.

The app they developed, VisiCalc, changed businesses forever and became the reason many people bought Apple IIs. "VisiCalc took 20 hours of work per week for some people, turned it out in 15 minutes, and let them become much more creative," Bricklin says.

The Apple II ended up selling fewer than 6 million units. The Apple II, and its competitors that would soon come, such as the IBM PC, did something important, though—they provided the scale the industry needed to start shrinking the components, making them much faster, and reducing their cost. That is what computing would need to do to sell not just a few million machines, but billions—which it did in the next two paradigms.

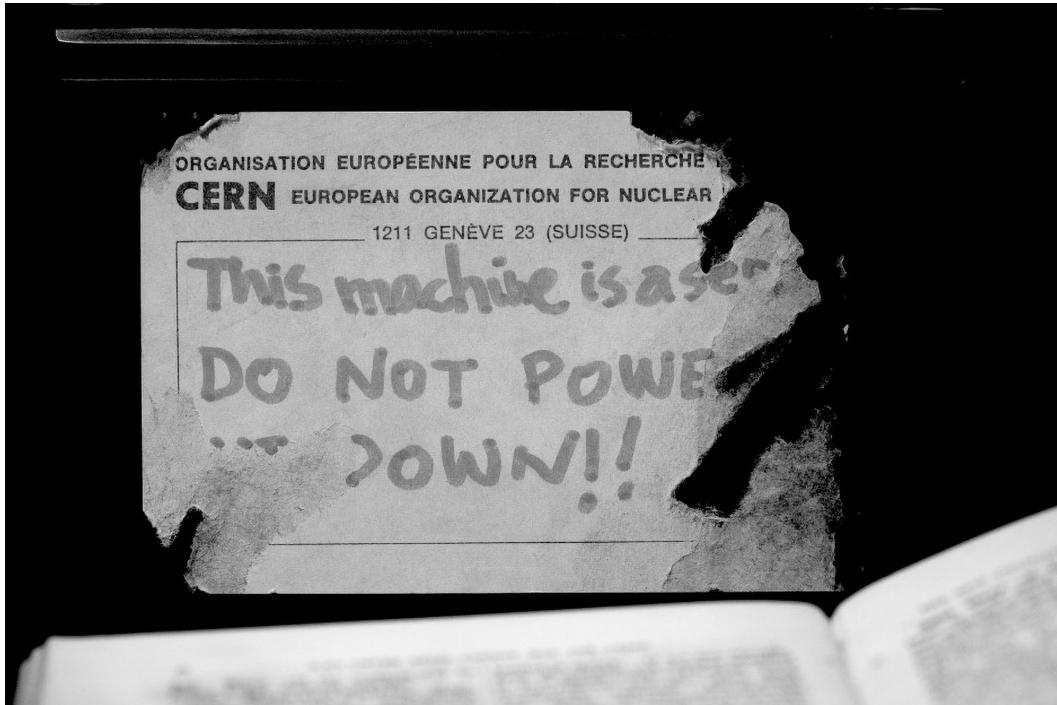


Figure 2.2: Label on Tim Berners-Lee's NeXT computer on which he invented the World Wide Web while working at CERN. The NeXT operating system still survives in today's Macintosh and iOS, but is really the best example of concepts invented at Xerox Palo Alto Research Center in the early 1980s. Photo credit: Robert Scoble.

Paradigm Two—Graphical Interfaces and Thinking

In 1984, the launch of the Apple Macintosh computer brought us graphical computing. Instead of typing text commands to, say, print a document, you clicked an icon. That made computing much easier but also enabled the launch of a new series of desktop publishing apps. By the time Windows 95 arrived, finally bringing the idea to the mass market, the entire technology stack was built around graphics. This greatly increased the size of tech companies and led to profitable new lines of software for Microsoft and Adobe, setting the stage for Spatial Computing.

This graphical approach was much easier to use than learning to type in commands to copy files and print. Now, you could just click on a printer icon or a save icon. This increased accessibility brought many new users into computing. The thing to take away here is computing, with this and each of the paradigm shifts that followed, made a massive move toward working more like humans do, and made handling computer tasks much easier. Spatial Computing will complete this move (Google's Tilt Brush in VR still uses many of the icons developed in this era to do things like choose brushes or save/delete files).

It was a massive increase in the number of computer users (many stores had long lines of people waiting to buy Windows 95) that gave Microsoft, in particular, the resources to invest in R&D labs that led directly to the development of HoloLens 25 years after Windows 95's huge release.

Also, it helped many graphic designers move away from working on typesetting machines and brought them into computing. This accelerated with the popularity of the web, which also saw its debut on Windows 95 and Macintosh. When Tim O'Reilly and Dale Dougherty popularized the term Web 2.0 in 1994, even Bill Gates didn't understand how important having people interacting on web pages would be.

Weblogs were springing up by the millions and e-commerce sites like Amazon and eBay were early adopters of techniques that let parts of web pages change without being completely refreshed. Today, WordPress is used by about 20 percent of the web but back then, Gates and his lieutenant Steven Sinofsky didn't see the business value in Web 2.0, refusing to consider a purchase after the coauthor of this book, Robert Scoble, suggested such a thing when he worked as a strategist at Microsoft.

The web was starting to "come alive" and desktops and laptops were too with new video gaming technology. Nvidia, born in the 1990s, was now seeing rapid growth as computing got cheaper and better. Where a megabyte of RAM cost \$450 in the late 1980s, by 2005, a gigabyte was running at \$120. While that was happening, internet speeds escaped the very slow modem age and were increasing to such speeds that new video services like YouTube and social networks, including LinkedIn, Twitter, and Facebook, were able to appear. Technology continues to get cheaper and smaller to this day, seemingly following Moore's Law. Today, 21 billion transistors fit into a chip the size of your fingernail (like the one that does self-driving in a Tesla) and memory now costs about \$15 for 64 GB chips. It is this decrease in cost and increase in capabilities that is bringing us Spatial Computing.

This new user interface paradigm, while easier than typing text commands, was still difficult to use for many. More than one phone call to Microsoft's support lines demonstrated that many had difficulty figuring out how a mouse worked, and even dragging icons around the screen was daunting and scary to many. Those of us who grew up around computers saw it as easy, but many didn't. While the visual metaphors were there, the disconnect between moving your hand on a mouse on a desk while controlling a cursor on the screen meant that computers still didn't work the way we did. While most figured it out, computing had another problem—computing wasn't designed to fit in your hand or pockets, which is the vision that Douglas Engelbart, among other pioneers, had for all of us. Engelbart was the genius who, back in the late 1960s, showed the world the technology that it wouldn't get until the Macintosh arrived in 1984. Before he died, he told us he had an unfinished dream: of making computing even more personal, to the point where you could communicate with computers with your hands and eyes. He predicted not only the move to mobile but also the move to truly Spatial Computing.

This brings us to paradigm three: Mobile.

Paradigm Three—Mobile

Humans aren't happy being tied to desks to do their work, and that enabled a new industry and a new paradigm of personal computing to arrive—one that brought computing off of desks and laps and into your hand. This paradigm shift enabled billions to get on the internet for the first time (we've seen very poor people in China and other places riding bikes while talking on their smartphones). This would be the platform that many new companies would build upon thanks to new sensors, ubiquitous data networks, and new kinds of software that were designed for these devices we all now hold.

This third technology shift started in places like Toronto (RIM Blackberry) and Helsinki (Nokia). For years, these two companies, along with Palm, with its Treo, and a few others, started a new technology industry direction. They produced products that fit in your hand and didn't seem to be very powerful computers at the time. Mostly, they were aimed at helping you take a few notes (Treo) or make a call while entertaining the ability to take a photo too (Nokia) or send a few text messages to coworkers (RIM's Blackberry). This turned into quite an important industry. Nokia alone, at its peak in 2000, accounted for four percent of its country's GDP and 70 percent of the Helsinki Stock Exchange market capital.



Figure 2.3: A Microsoft Windows Phone, circa 2006, sits next to the original mouse, circa 1968, on Douglas Engelbart's coffee table. Photo credit: Robert Scoble.

What they didn't count on was that Steve Jobs would return to Apple and with the help of Microsoft, who poured capital into the failing Apple of the late 1990s, bring Apple back, first by rejuvenating the Macintosh line to keep its faithful happy, then with the introduction of the iPod. Now, the iPod didn't seem to be a device that would help Apple compete with Blackberrys, Treos, and Nokias, but it helped Jobs build teams that could build small, hand-held devices and figure out how to market them. That effort crushed other portable audio players and did so in such a way to give Jobs the confidence, and cash, to invest in other devices like phones.

A few of the engineers who built the iPod were asked by Jobs to come up with ideas for a new device that would merge what they learned with the iPod with added phone capabilities. The early prototypes looked much more like an iPod than the product we all know today.

That team all carried Treos and studied the other devices, Andy Grignon told us. He was one of the dozen who worked on the first prototypes. They saw that while early devices were useful because they could be carried around, they were hard to use for most things other than making phone calls. Many didn't have keyboards that were easy to type on, for instance. Even the ones that did have keyboards, like the RIM devices, were hard to use to surf the web or do things like edit photos.

He told us that Jobs forbade him from hiring anyone who had worked on one of these competitive products, or even from hiring anyone who had worked in the telecom industry. Jobs wanted new thinking.

On January 9, 2007, Steve Jobs introduced the iPhone. That day, we were at the big Consumer Electronics Show getting reactions from Blackberry and Nokia execs. They demonstrated the hubris that often comes from being on top: “Cupertino doesn’t know how to build phones,” one told us. They totally missed that there was an unserved market—one that not only wanted to use devices while walking around but also wanted to do far more than just make a call or take a photo once in a while. Their devices were too hard to use for other tasks, and their arrogance kept them from coming up with a device that made it easy.

Around that time, the web had become something that everyone was using for all sorts of things that Tim Berners-Lee, the inventor of the web, could never imagine. With iPhones and, later, Android phones, we could easily use the full web on our mobile devices while walking around, all by using our fingers to zoom into articles in the New York Times, just like Steve Jobs had demoed months earlier from a stage in San Francisco.

It was this combination of an easy-to-use device, along with sensors, that could start adding location-based context to apps that formed the basis of many new companies, from Uber to Instagram, that were born within years of the iPhone launch. This showed something significant had happened to the world of technology and set up the conditions for the battle over where the tech industry will go next: Spatial Computing.

Paradigm Four—Spatial Computing

You might notice a theme here. Each paradigm builds upon the paradigm that came before, bringing real breakthroughs in user experience. With our mobile phones, tablets, and computers, there’s still one glaring problem—they don’t work like humans do. Paradigm four is bringing the perfect storm of all usability breakthroughs.

Even a young child knows how to pick up a cup and put it in the dishwasher or fill it with milk. But this same child is forced to use computing that doesn’t work like that. Instead of grabbing with their hand, they have to touch a screen or use a mouse to manipulate objects on a flat screen.

In Spatial Computing, they will just grab a virtual cup like they would a real one in the real world. This will enable many new people to get into computing, make all of our tasks easier, and introduce new ways of living.

This move to a 3D world isn't just happening in computing, either. We experienced an off-Broadway play, "Sleep No More." It was a remake of Shakespeare's Macbeth, where you walk through sets, with action happening all around. Even New York plays are starting to move from something confined to a rectangular stage to one that surrounds us in 360 degrees. It's a powerful move—one that totally changes entertainment and the audience's expectations of it.



Figure 2.4: A man wearing an Apple Vision Pro taking advantage of its Spatial Computing capabilities. Photo credit: Apple.

If Spatial Computing only introduced new 3D thinking, that would be absolutely huge. But it's joined by new kinds of voice interfaces. Those of us who have an Amazon Echo or a Google Home device already know that you can talk to computers now and ask them to do things. We are now able to have entire conversations with computers. Also, there are powerful new AI systems that can "see" things in your environment. Computer vision is making getting information about things, plants, and people much easier.

A perfect storm is arriving—one that will make computing more personal, easier to use, and more powerful. This new form of computing will disrupt seven industries, at a minimum, which we will go into deeply in the rest of this book. In the next section, we'll look at how this storm of change is to impact technology itself.

Six Technologies

The storm will bring major advances in six new technologies: Optics and Displays, Wireless and Communications, Control Mechanisms (Voice and Hands), Sensors and Mapping, Compute Architectures (new kinds of Cloud Computing, for instance), and AI (Decision Systems). For the rest of this chapter, we'll explore the implications of the fourth computing paradigm for each of these technologies.

Optics and Displays

The purest forms of Spatial Computing will be experienced while wearing a head-mounted display personified in the form of glasses. In these glasses, you will experience either Augmented Reality (where you see some of the real world along with some virtual items either placed on top, or replacing real-world items) or Virtual Reality (where you experience only a virtual world).

Paul Milgram, an industrial engineering professor who specializes in human factors, went further into depth than just dividing the world into these two poles. He said that these two were part of a reality-virtuality continuum, which included Mixed Reality (where both the real and virtual are mixed). People will continue to mix up all the terms, even though Apple has decided to call it Spatial Computing.



Figure 2.5: DigiLens prototypes, as displayed at SPIE Photonics West 2019, a conference focusing on optics. Photo credit: Robert Scoble.

Other terms used to describe various experiences on the spectrum between augmented and Virtual Reality include computer-mediated reality, extended reality, simulated reality, and transreality gaming. Now, you probably get why we prefer to just lump all of these under one term: Spatial Computing.

There is already a range of Spatial Computing products on the market, from the Apple Vision Pro on one side to Microsoft HoloLens on another, with an approach that puts virtual images on top of the real world in a very compelling way. The Apple Vision Pro does suffer from its heavy weight and low battery life, but its imaging is dazzling. It's also really expensive at \$3,500, which does not allow for mass consumption. We hear that there will be a newer and less expensive model (costing something like \$1,500) that will be coming within the next two years. In talking with Qualcomm, which makes the chipsets in most of these products, we learned that there are two approaches that will soon dominate the market: one where you see through the virtual layer to the real world, and one that passes the real world to your eyes through cameras and screens.

Both fit into the Spatial Computing family since they present computing you can move around in, but they have completely different use cases.

Examples of see-through devices are Magic Leap 2 or Microsoft's HoloLens, which let you see the real world through the lenses a wearer sees through, but use a Wave Guide to display virtual items on top of the real world.

In November 2019, it was leaked that Apple would bring HMDs that bring something else: a completely opaque display that, when worn, presents the real world "passed through" cameras to its tiny little high-resolution screens. This passthrough approach, first introduced in the spring of 2019 in the Varjo VR-1, brings you a super-wide **field of view (FOV)**. That's something that the Magic Leap 2 and HoloLens can't yet match, along with extremely high resolution and better color than Magic Leap 2 and HoloLens can bring you. The Meta Quest 1 and 2 had limited pass-through capability (bad resolution and monochrome). It was only with the Meta Quest 3 that passthrough was much improved with full-color high-resolution technology. Apple came out with its passthrough Apple Vision Pro, which trumps the Meta Quest 3 in terms of technology and overall quality. Apple glasses (see-through, not passthrough devices) are not expected until around 2028, with Meta probably beating them to the punch.

The worry that people would not be able to see your eyes with passthrough devices was solved by Apple with its Vision Pro, even though the eyes that people see are a very good, rendered representation of a person's eyes. So, in this way, people can see you are available to talk with.

However, the Vision Pro won't be good for the military or police, who need to see the real world in order to shoot guns. The latency, no matter how small, and imperfections in bringing the real world to your eyes in such an approach might prove troublesome in applications that require immediacy. Similarly, we wouldn't trust them in everyday situations where we could be in danger if our awareness is compromised, like workers using a bandsaw, someone driving a car, or even people trying to walk around a shopping mall.

That said, a passthrough device would be amazing for replacing monitors. Another way to look at it is that passthrough will be biased toward Virtual Reality, where you see mostly virtual worlds or items and very little, if any, of the real world. See-through displays, on the other hand, will be biased toward Augmented Reality, where you mostly see the real world with virtual characters or items, or virtual screens overlaid or replacing parts of the real world. Each is a member of the Spatial Computing family and has different use cases and benefits.

We don't know which approach will win with consumers, but we have come up with some theories by trying out the different approaches for ourselves.

Enterprise workers who need to walk around machines, or police/military users, will tend to go with something like the HoloLens.

These devices use optics systems that guide light to your eye through glass with tiny structures, or mirrors, that bounce light projected from tiny projectors or lasers into your eyes, giving you an image. Why do we say workers or the military will go with these? Because they let you see through the glass to the real world. The analog world is hard to properly represent virtually with low latency, and these optics don't even try.

The devices that require lots of cameras, sensors, and powerful GPUs to properly align virtual images in the analog world will remain fairly big and heavy. The HoloLens weighs more than a pound and it doesn't even include a 5G radio or batteries that can keep the thing powered for long enough for it to be used for more than a few hours.

We can't imagine running, or skiing, with such a big contraption on our heads, not to mention that they make a wearer look antisocial if worn out in public. This explains why there is a small, if not underpowered, device like Meta's Ray-Ban Wayfarer. We find these smaller devices lacking in the ability to present text in a satisfying way and they have many limitations, from having a very small field of view to not having enough power to do Augmented Reality, which is where we believe the real power of on-face wearable devices lies. These are great for wearing while doing, say, inventory at a warehouse, or seeing notifications like a smartwatch might display, but they can't display huge virtual monitors and can't even attempt to use virtual or true Augmented Reality-style entertainment apps.

Bigger devices with optics from companies like Kopin or Lumus, which makes the display in Lenovo's glasses, and DigiLens, which got an investment from Niantic, the company behind Pok  mon Go, are pushing the bleeding edge, with many others coming soon.

You should be aware of these different approaches, though, and the pros and cons of each. For instance, people who mostly work at desks, whether sitting at a Starbucks or in a corporate office, won't often need to see the real world, so they may prefer the wider field of view, better color, and higher resolution of passthrough devices like the Varjo, and what we expect Apple to be working on, among others.

Why do we see passthroughs as the ones that may win with consumers by 2026? New micro-LED displays are coming to the market from companies like Sony and Mojo Vision, which have super high-resolution displays that are smaller than your fingernails.

These displays offer better power utilization and some can do things like completely turn the sky you are looking at dark, or replace it with virtual images. That lets wearers experience both VR and AR, often one after another, in all sorts of lighting conditions, potentially with wrap-around screens, at a lower cost than the techniques used in HoloLens.

We can see a world where we'll own two or maybe even three pairs of glasses for different things. Additionally, we might use a very different headset, like the Varjo, when we go to a shopping mall to play high-end VR in Location-based Entertainment retail settings like you see at Sandbox VR. There, you pay \$20-\$30 or more to have much more immersive experiences than you can at home. However, for the purposes of this book, we'll focus on enterprise and consumer uses of Spatial Computing and not these high-end uses—although you should try them out since our families and team members love them as well.

One thing we want you to realize is that much better devices will soon arrive—some may have even been announced between us finishing writing this book and you reading it. Expect this section of our book to be out of date quickly after printing as technologies come out of the lab.

As we've been talking with people about these different approaches, many people believe that small GPUs will not be able to power such high-resolution displays. That might be true if displays behaved like the ones on your desktop, which have the same resolution all the way across the display.

VR pioneer Tom Furness told us that won't be true soon, thanks to a new technique called "foveated rendering." What he and his researchers realized, is that the human eye can only perceive color and high resolution in the middle of the eye, the fovea. Since Spatial Computing devices will be worn and some include eye sensors, they can play visual tricks with our eyes, he says, by only putting high resolutions, or tons of polygons, right in the center where our eyes are looking.

The outside boundary of the screens could be lower resolutions, or, at a minimum, display fewer polygons. This would lead to “fooling” your eyes that the entire display is, say, an 8K one, while it really is only the center of the display that is that high a resolution. Already, this technique is being used in some VR headsets, like the Varjo, and products that use this technique can double apparent resolution without putting extra strain on the GPU or computers driving them.

Other problems also need to be solved. With the HoloLens, you can’t bring virtual items close to your face. Why is that? Because of two terms: accommodation, which is how the muscles in your eyes react to closer-focus items, and vergence, which is how your eyes move toward each other when focusing on something close. Those two actions are hard for current optics and microscreens to support.

The holy grail of displays was actually invented about 20 years ago by Furness. It uses a low-power laser to scan a single pixel across the back of your retina. No one has been able to manufacture such a monitor at scale yet. However, Furness said prototypes in his lab actually helped a few people to add sight to their blind eyes since the laser penetrates scar tissue that some people built up due to birth defects or injuries to the back of their eye. We hope to see a working version of such a monitor, which, if a company can build it, promises to use very little power and give amazing images due to being painted right onto our retinas.

Now that we’ve discussed the future of Optics and Displays, let’s look at the other pieces of the Spatial Computing puzzle that are also rapidly evolving.

Wireless and Communications

The screens mentioned in the previous section are going to be data-hungry. If our sources are correct, the wearable screens on our faces will need 8K and probably much higher resolution videos to give you a sharp view of the world around you. We are hearing of companies testing glasses with 32K videos that wrap around you in 360 degrees.

For such a viewing experience to be true, you’ll need a huge amount of data—all while the glasses are also uploading data streaming from cameras and sensors up front.

Plus, engineers in labs are telling us that they are trying to remove as much computing from the glasses as possible and put that computation up on new kinds of cloud servers. For such a scheme to work, we’ll need not only much more bandwidth than we have today but also very low latency.

Enter 5G. 5G brings us three benefits. First, it gives a very high bandwidth. Users are getting about 10 times more bandwidth than they used to get with LTE on mobile phones. Second, it provides low latency. 5G has a latency of about two milliseconds to a cell tower.

The current latency is about 10x slower. That matters if you are trying to play a video game with someone else. Imagine everything lags or is slow. The third benefit comes to bear if you visit sporting/concert arenas with thousands of other people. 5G promises many times more devices can connect per antenna. We remember the day we went to the World Series at AT&T Park in San Francisco and couldn't send a text message or make a phone call on our iPhones, despite AT&T putting the best radios it had into the park.

As of 2024, many of you have 5G and it has already been put into sports arenas, and cities with high population densities. 5G also has some significant problems. The highest frequencies can't go through walls easily and you need to be much closer to antennas than with LTE. Telecom experts tell us that rollouts have been slow in some neighborhoods and you will probably need to upgrade your WiFi in order to distribute such fast bandwidth inside office buildings and homes.

When will everyone have 5G in their home? If you live in a high-density city, probably by year-end 2024. If you live in a rural area, though, you might not see it for most of the 2020s due to it being necessary to be very close to a 5G tower for true 5G. That said, in 2019, we got an AT&T fiber line and WiFi 6 in our homes, which already gives us most of the benefits of 5G without waiting. WiFi 6 is the newest version of the WiFi standard and shipped on iPhones for the first time in late 2019. Unfortunately, very few people have access to a gigabit fiber line. Our WiFi 6 router cost about \$400, and we bought another \$500 worth of extenders. Very few people will want to put more than \$1,000 into new WiFi 6 equipment just so they can have the fastest speeds in the neighborhood!

Another place where 5G has been rolled out more quickly? Inside factories. Here, 5G pays real dividends because virtualized factory floors with virtual twins of machinery can make things much more efficient, and many factories have thousands of sensors and robots that can overwhelm older communications infrastructure.

It is the new software architectures that 5G enables, however, that are most exciting. As discussed in *Chapter 1*, toward the end of 2024, Spatial Computing will be getting a big boost from release 18. This will make playing games and seeing other kinds of visuals in headsets, such as the Apple Vision Pro, a really astounding thing. Imagine creating a video with Generative AI and then being able to seamlessly view it with hardly any latency.

The thing with all this bandwidth is that it's very hard to use on older equipment. Even with tons of 4K video streaming, we can't slow down our gigabit internet line. It's Spatial Computing that will use all that bandwidth that 5G, or a WiFi 6/Gigabit internet line, will bring. To control that we'll need a much more efficient way to control the multiple video screens and the Augmented Reality that will surround us as we wear new Spatial Computing glasses.

This is why Qualcomm is already building wearable devices that have seven cameras, most of which will be used for new kinds of controls for this new highly interactive world.

Control Mechanisms (Voice, Eyes, and Hands)

When we put an Apple Vision Pro on someone for the first time, what do they try to do? Grab one of the visuals they see with their hands. Obviously, this is something that we would never attempt to do with a 2D interface such as a computer screen or a smartphone screen. And the Apple Vision Pro has really good hand gesturing, along with eye gaze tracking (voice not so good yet). Even Meta is considering not selling the future versions of the Quest with hand controllers!

This human need to try to touch, grab, and manipulate is so powerful that it's easy to see that controlling computing inside an AR or VR headset won't be anything like it is on a traditional laptop or desktop machine with its mouse and keyboard. There are so many business and consumer uses in Spatial Computing for this direct kind of manipulation, that is, using your hands.



Figure 2.6: A woman wearing an Apple Vision Pro uses her hands to gesture. Photo credit: Apple.

At the same time, new voice-first technologies are coming fast and furious. Apple has been working on a new Siri, which we expect could run the entire operating system of a Macintosh simply by talking. It has already taken a good step in integrating Siri with OpenAI's ChatGPT, announced during Apple's 2024 Worldwide Developers Conference (WWDC)¹. Many of us are already used to using voice for some things, either with Siri on iPhones; Alexa on Amazon Echo devices; or Google's Assistant on its Android phones, apps, or Google Home devices.

And devices like Limitless, which uses Generative AI to take notes or do other work, are the wave of the future.

While most people will judge devices by how well they handle hand gestures, we see that it's the combination of input methods, whether it's using your voice, your hands, your eyes, or even typing or moving a phone, that makes Spatial Computing so powerful.

It's the first time we can really experience hands-free computing. Surgeons, for instance, are starting to use Spatial Computing because they can control the cameras and other devices by talking to their devices.

Haptic controls are also evolving quickly, which adds realism to many use cases. The standard controllers that come with, say, an Oculus Quest or HTC Vive, do some, like shaking when you shoot a gun. However, there's an entire industry springing up as Spatial Computing glasses and VR headsets become more popular, from companies that make gloves that let you touch and feel things, to suits that let you capture your entire body motion, and even, in the case of the \$15,000 Tesla suit, let a wearer experience feeling hot or cold, along with other sensations passed to your skin through dozens of transducers.

That said, the standard sensors in headsets alone bring major new capabilities, even if you don't add on some of these other accessories.

What is really going on is that we're seeing a convergence of technologies that are making all of this possible. Sensors are getting cheap enough to be included in consumer-grade devices. AI is getting good enough that we are close to having full conversations with our computers, as Google first demonstrated in 2018 with its Duplex demo. There, an AI system called a local business and talked with the human answering the phone and the human was none the wiser.

Computer vision and machine learning make using hand gestures much more accurate. Andy Wilson showed us how gestures worked at Microsoft Research way back in 2005 (he directs the perception and interaction research group there; you can learn more at <https://www.microsoft.com/en-us/research/people/awilson/>). He wrote algorithms that could "see" his fingers with a camera. Another algorithm was taught how to tell if he touched his finger and thumb together. Then, he zoomed in on a map just by using his hands in mid-air. Today, we see that same technique used in Microsoft's HoloLens, even though it has become much more advanced than the simple touch-your-fingers-together algorithm.

Back then, though, gestures only did one or two things and weren't accurate if the sensor couldn't clearly see your thumb touching your finger. Today, AI systems pick up small patterns on the back of your hand from your muscles contracting so that many more gestures work, and the ones you try might not even be fully visible to the sensors on the front of your Spatial Computing glasses.

These capabilities make Spatial Computing devices awesome for remote assistance, hands-free work, and all sorts of new interactions with computing or robots. Snap your fingers and you can be shooting at monsters crawling toward you in a video game.

The thing businesses need to know here is that users will come at this world from where they are, and you will need to do a lot more testing with people to make sure they can use your services well. Accessibility is a huge opportunity. Apple demonstrated this when a man who had no use of his hands, working in a wheelchair, could still send messages, manipulate photos, and perform other tasks, all with his voice.

It is hard to imagine just how deeply user interfaces will soon change. Eye sensors will facilitate some major leaps in the way we interact with our technology. We can see a day when we are in the shopping mall and ask our glasses "Hey, Siri, what's the price of that?" totally expecting an answer about the shirt we are actually looking at. Indeed, all the major companies have purchased eye sensor-producing companies. In part, this is because eye sensors can be used to further develop high-resolution monitors that pack tons of polygons into the spot where you actually are looking using the foveated rendering technique. However, we bet they see the business opportunities of making user interfaces understand us much better by knowing what we are looking at, along with the biometric security that will come, too.

Another need for eye sensors is biometric-multifactor identity. It turns out our eyes are like our fingerprints: uniquely ours. So, if our glasses are put on by someone else, they won't be able to access our private information until they hand them back to us. Add that to identity systems that use voice, heart rate, and the blood vessels in the skin that your cameras can see, where your human eye can't, along with patterns like gait and hand movements, and computing will soon be able to know it's you at a very high degree of accuracy, increasing the security of everything you do and finally getting rid of passwords everywhere.

Eye sensors will also be loved by marketers since it's possible they will know if you actually looked at an ad or got excited by a new product display at a store. These new capabilities will lead to many new privacy and control concerns.

Speaking of Siri, as of 2024, it isn't all that good. However, Apple is working on improving it to the point where you can do everything, from finding a restaurant to editing a video, just by talking to it by utilizing its integration with ChatGPT.

Foursquare actually has an answer to Siri's current disappointing performance, along with an API so other programmers can get to it. Siri, though, was hard-coded and Apple's programmers haven't yet gotten around to such an obscure feature, so it fails by bringing in an answer from Bing that makes no sense.

Amazon Echo/Alexa and Google Assistant/Home demonstrate something is going on here: we are increasingly using our voice to control our computing. Much better AI is coming between now and 2027 that will adapt to even difficult-to-understand voices, as well as much more complex queries that go far beyond the Foursquare example.

As systems get more data about where we are, what we are doing, what we are looking at, what or who we are interacting with, and what our preferences are, virtual beings and assistants will become much better. Businesses will increasingly run on them, and customers will choose businesses that better support our assistants. A "Hey Siri, can you get me some Chinese food for lunch?" type of query might soon decide that your Chinese restaurant is the one it'll bring thousands of customers to. Why you, rather than a competitor? Maybe you put your menu online, maybe you are better at social media, maybe you have better Yelp ratings, or maybe your staff is wearing Apple glasses and can answer queries for Apple customers faster so Siri starts to prefer your business.

Put all of these together and you'll see a computing system that works the way you do with other people; by gesturing and touching, talking, and using your eyes to focus attention on things.

That said, a major way we'll navigate around is, well, by moving around, and there's some major new technology that's been developed for sensing where we are and what surrounds us. You'll see versions of this technology used in robots, autonomous cars, and Augmented Reality glasses.

Sensors and Mapping

We humans take moving around the real world for granted. It seems easy to us, something we've done ever since we were born after our parents bundled us up and brought us home.

Computers, though, couldn't understand this real world until recently. They didn't have "eyes" with which to see it and didn't have even a basic understanding of being able to move around, that is, until the mobile phone came along and we started teaching it by building digital maps.

At first, those digital maps were mere lines where streets were, with a few dots of data along the street. Today, however, our glasses, cars, and robots in the street are mapping that street out with billions of points of data. Now, this same process is happening inside our factories, hospitals, shopping malls, and homes.



Figure 2.7: Here, a street is mapped out in 3D by Luminar Technologies' solid-state LiDAR.

Photo credit: Luminar Technologies.

Mapbox, which is the map provider that hundreds of thousands of apps use, showed us the future of maps that was only possible once mobile phones became ubiquitous enough to get huge amounts of data on every street. It showed us how they could build an entire map of a freeway, including lanes, via huge numbers of mobile check-ins. Every time you open Yelp or Foursquare, or use Snapchat, among other apps, it gets another little point of location data. If you get enough of these points, which they call “pings,” a new kind of map is possible due to mega-huge databases of all that data, combined with machine learning to make sense of all that data streaming into its servers.

Then, there is data collected by cameras or new kinds of 3D sensors, like time-of-flight lasers, which measure the world out very accurately by figuring out how long a beam of light takes to reflect off of a surface and get back to it. Your 3D sensor in a modern iPhone works like that. Today, Apple mostly uses that sensor for Augmented Reality avatars or face detection to unlock your phone, but in the future, those sensors will do a lot more.

A Tesla car, driving down the same street, images the street as a point cloud of 3D data, and there are many streets where a Tesla drives down the street every few seconds, gathering even more data.

Soon, every stop sign, tree, pole, lane marker, restaurant sign, and far more insignificant details will be mapped by many companies. Apple alone drove four million miles with 360-degree cameras and 3D sensors in an attempt to make its maps better. Other companies have driven even further.

These early efforts in mapping the world are great, but the data points that all these companies have collected are going up exponentially. The next frontier will be to map out everything.

Our friends who have worked at Uber tell us it is keeping track of the sensor data coming from the mobile phones of its drivers. It can tell if streets have a lot of potholes just because those driver's phones are shaking. The auto insurance company, Go, says it has maps of which streets have more accidents than others. It and other companies are laying all sorts of things on top of the digital map, from sensor readings coming off of machines to other location data from various databases, similar to how Zillow shows you the quality of school districts on top of homes you might be considering buying.

Luminar Technologies' CEO, Austin Russell (he runs a company that builds LIDAR sensors for autonomous cars), told us how his sensors work—they map out the real world with solid-state sensors that can see up to 250 meters away. The high-resolution point cloud that is generated is processed by AI that quickly identifies literally every feature on the street. Other companies take that data and build high-resolution contextual maps that future cars can then use. Russell explained to us how these maps are nothing like maps humans have ever used before. These maps, he says, aren't a simple line like you would see on your mobile phone, but a complete 3D representation of the street, or surroundings of where you are.

The glasses of the future will do something very similar. Companies like 6D.ai, which was acquired by Niantic, had been using the cameras on mobile phones to build a crude (as of 2019) 3D map of the world. 6D.ai had been calling it an AR Cloud, but what it really is is a sheet of polygons stretched on every surface that a particular camera sees. AR Cloud is a new term that incorporates a technology called SLAM, which stands for Simultaneous Location and Mapping. That's the technique developed by NASA to have the Mars Rover navigate around the surface of Mars without humans being involved. Basically, it is how a robot can see the real world and move around in it. AR Cloud takes that further by giving a framework for where virtual items can be placed.

This is how Augmented Reality works at a base level. In order to put a virtual item into the real world, you need a “digital twin” of that world. We don't like that term, because in reality, we see some places having dozens of digital twins. Times Square in New York, for instance, will be scanned by Apple, Facebook, and Microsoft, along with dozens of car companies and transportation companies, and eventually robots from Amazon and other companies, which will be rolling through there delivering food and products. Each builds a digital twin.

This digital twin need not look anything like the real world to humans. At Niantic, the folks who make *Pokémon GO* have such a digital twin with only data of where people play games built on top of its platform—parks, shopping malls, sports arenas, that kind of thing. The data in its platform doesn’t look like the real world to a human. Their “map” of the real world then converts that data into contextual information about the real world. Is there water nearby? Is this a park where children play, or adults hike? That data lets them build better games, by putting specific *Pokémon* characters near water, and others, where there are, say, tons of forests. Their worldwide database is doubling in size every few months.

The holy grail, though, is to have a detailed map of every surface that surrounds a human. Already, at factories like the ones Audi or Volkswagen run, they have done just that. They used cameras and 3D sensors to build a high-resolution virtual version of their factory floors using a system from a company called Visualix, which tracks everything in that factory down to a centimeter accuracy, its former CTO Michael Bucko told us. Why? For two reasons: training and design.

It can use that virtual factory floor to train employees on new jobs using VR headsets before they even see the real factory. It can also use that same virtual factory floor to see how to redesign part of it to hold a new set of robots, for instance, and design workspaces for the humans that will need to interact with the robots. It won’t stop there, either. Workers, who wear Augmented Reality headsets, will be able to be assisted during their work while on the real factory floor. Re’flekt already built such a system for Audi workers. They can even leave videos on the real factory floor for the next shift to explain how to use a new piece of equipment. This is especially useful if a worker is retiring and will be walking out the door with dozens of years of knowledge about how said equipment actually works. These videos can be left on that real equipment by attaching it to the “digital twin” that is actually in a database up in the cloud. To the human, it looks real, but it is all virtual and enabled by sensors and cameras that have mapped out the factory floor.

Bucko says its hundreds of customers were using Visualix’s SLAM-mapping and positioning technology to build the foundation for Augmented Reality, as well as the digital twin that thousands of robots will use to navigate inside, say, a huge warehouse or retail store.

Your kitchen will soon be the same as that factory. The cameras on the front of your glasses will build a 3D map of your kitchen (Microsoft’s HoloLens and even the \$500 Oculus Quest 3 have four cameras that are constantly mapping out the room you are in, at a high resolution). It won’t take your privacy away—that’s how it works, so that you can play Virtual Reality or use Augmented Reality. On the HoloLens, sometimes, you even catch a glimpse of the polygon sheet that it’s producing on your walls, tables, and floor. It then uses that sheet of billions of little triangles to put virtual items on top.

Today, we don't have enough bandwidth or power in our devices to upload that sheet to the cloud, but we can see how, in the future, you will want every room in your home mapped out this way, along with every place you typically use these devices in: your car, your office, even stores. Today, these maps, or sheets of polygons, are stored locally, partly for privacy reasons, but also partly because, again, the small batteries and processors inside just can't handle a lot of uploading and downloading and comparing of these sheets yet. That will start to change dramatically this year, as this process of seeing the real world and uploading and downloading a digital twin of it gets more efficient. Visual mapping with cameras, LIDARs, and 3D sensors isn't all that's going on here, either.

Once these sheets of polygons, or even the higher-resolution point clouds, hit a computer, you'll see systems that categorize this data. MobileEye, at CES 2019, showed what we mean. This autonomous car company, started in Israel and now owned by Intel, not only builds a digital twin of the real world but also, as users drive by, say, a restaurant, ingests that data into its AI systems and then categorizes that data into a useful database that sits on top of the map. Now, it knows where your favorite Chinese restaurant is. It also knows what is next to it. So, future drivers can now say "Navigate me to my favorite Chinese restaurant" and the autonomous car will take them right there.

The dream is to go further than any of this—to a memory aid. After all, if these imaging and categorization systems are already able to map out where your favorite restaurants are, they can also do things like remember where you left your phone or your keys, even warning you if you walk away from a table in said restaurant without picking up your phone. They could say to you, "Hey dummy, you left your phone on the table inside the restaurant."

Soon, these databases will remember everything in your world and notice if you move things around, or where you left your keys. Why? The demand on these companies will be to build much more realistic augmented and Virtual Reality experiences. Virtual beings will even sit on chairs with you in your kitchen and have conversations with you. To do that properly, it will need to be fully situationally aware—not just know there's a chair there, but the context of a chair. After all, don't you behave differently when sitting in a chair at a church, school, or noisy bar? So, too, will the virtual beings in our future glasses.

That sensor and map data will lead to new computing architectures, too, particularly when paired with the high bandwidth of 5G and new WiFi.

Computing Architectures

Computing architectures have also gone through major paradigm shifts, and as we move into Spatial Computing, will see another major shift: to decentralized systems. Databases that once used to be stored on a single server, then multiple ones, moved to virtual cloud servers. We saw that move first-hand working inside Microsoft and particularly Rackspace. When Rackspace started in the late 1990s, it installed servers in racks for customers. Today, those customers buy virtual servers that can be started up in far less than a second.

This move led to Amazon Web Services being the dominant way that start-ups build their infrastructure. Today, Amazon offers dozens of different kinds of virtualized servers, from databases to ones with beefy GPUs to do image processing on. It also started out as having one big data center, moving to having many around the world in regions.

The goal is to get computing as close to users as possible. Some of that is for redundancy, yes, so if a data center in one region fails, the others kick in and the service stays up. Today, though, as we move into Spatial Computing, there are another couple of reasons—to reduce latency and the costs of building systems that will see very heavy workloads. See, we are still controlled by physics. A packet of data can't travel faster than the speed of light, and, actually, it is slowed down by each piece of equipment it travels through.

Plus, as we move our businesses from, say, things that work like e-commerce systems that serve web pages, to ones that look and work more like interactive video games, traditional infrastructure that's centralized will start to crack and break, not to mention becoming very expensive.

This is why most companies no longer use Oracle to house their databases, preferring instead to house their databases on hundreds of virtualized servers, or even using serverless architectures that let companies like Amazon or Microsoft handle all that infrastructure work.

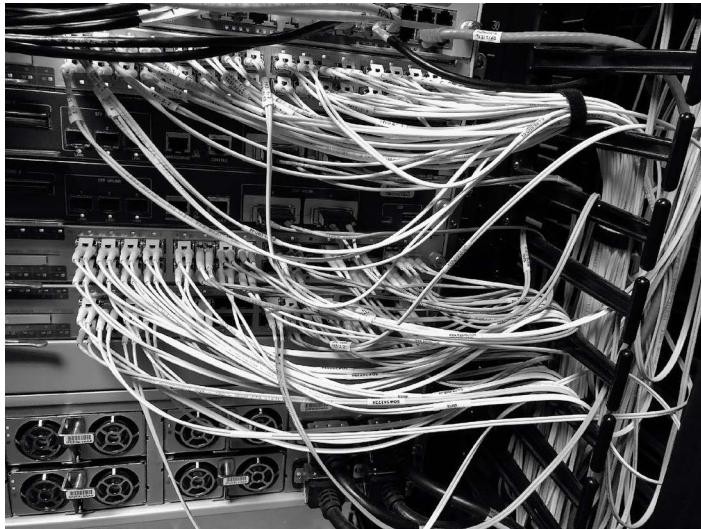


Figure 2.8: A bunch of wires serving 5G radios in Las Vegas' T-Mobile Arena. This is just a very small part of how Spatial Computing experiences will get to your glasses and mobile devices.

Photo credit: Robert Scoble.

As we move more fully to 5G, which gets that packet from a cell tower to your glasses or your phone with very little latency, the latency elsewhere in the system will become the bottleneck. Already, we are seeing new computing architectures that push the systems that draw polygons on your glasses' screen out to the edge of the network—in other words, on the glasses themselves or a box that's sitting in your home or office or somewhere else nearby.

But some data still needs to travel elsewhere. Think about if you are playing a virtual shooting game in VR. Now, if you are playing against a neighbor who lives on your street or in your apartment building, you might see everything slow down due to latency. Why? Well, that packet might need to travel from your home in, say, Mumbai, all the way to one of the mega-huge data centers in Las Vegas and back, passing through dozens of switches and routers, not to mention fiber or, worse, copper cable, to make it back and forth. Just because you have two-millisecond latency times to your cell tower won't mean that you have a good experience shooting at your neighbor in a video game, or collaborating with a coworker on a new product design.

As a result, businesses will increasingly be forced to buy cloud computing resources closer to the people playing the game. The standard architecture used by businesses is to spread data over three to ten regions, or data centers. That won't be enough, though, in this new highly interactive 5G-enabled Spatial Computing world. Packets getting to major data centers will still introduce too much latency and cost, because if everything is centralized, you need massive server power to deal with the loads that will soon come due to new customer demands like interactive VR experiences.

A new architecture is evolving that some are calling "fog." A new three-tier approach "cloud" for massive data centers far away, "fog" is for servers closer to you, and "edge" is for servers that are built into your devices or are very close to you in your homes or offices. Your main servers might be on Amazon or Microsoft Azure hundreds or thousands of miles away, which we still will call cloud, but a new set of smaller data centers will evolve that will serve parts of cities or neighborhoods. This new layer, since it sits between the huge data centers that Amazon, Microsoft, IBM, and Google run for their cloud infrastructure and users, is called fog. It's sort of between a real-world cloud (the huge data center) and the ground (for example, your computer), just like real fog is. In this way, many packets won't travel far at all, especially if you are only trying to shoot your neighbor virtually. These new smaller servers, or fog servers, don't need to be massive machines because they will support far fewer users. So, this three-tier approach will scale to many new application types.

Businesses that come into Spatial Computing will need to build using this new three-tier approach, pushing as much computing power out as close to users as possible. Older businesses will find that moving from a heavily centralized infrastructure where everything lives on huge data centers to one that's decentralized will be culturally and financially challenging, but it must be done to support not just these new use cases but also the many new users we expect (and the increasing demands on infrastructure we expect they will bring).

While that is all happening, we foresee a new set of infrastructures needed—content management systems. Designing standard business web pages that keep track of the images and text that need to be translated into local languages is fairly simple compared to having, say, different Pokémon characters for each city around the world. We visited EchoAR to get a look at one of these new Spatial Computing management systems. They showed us how users could even change their avatars and that the system would distribute those new avatars to other gamers in real time, or even corporate users around the world. And now, with Generative AI, assets could be generated on the fly, though generated 3D assets and video still have a while to go before they could be truly useful.

Add it all up and CIOs inside businesses will need to learn new systems and budget differently for the new workloads that are rapidly evolving.

AI (Decision Systems)

AI has gone way beyond telling the difference between dogs and cats. Today's AI, in the form of machine learning and deep learning, and often with the aid of Augmented Reality, is used for everything from building virtual cities to driving autonomous vehicles and robots to helping answer your emails.

All of these elements provide the situational "awareness" that is necessary for software running in a machine to present the appropriate images, and possibly accompanying sound to the person who is viewing Virtual Reality or Augmented Reality experiences or riding in an autonomous vehicle.

Before we go deeper into how AI is generally used for Spatial Computing, let's define what AI, machine learning, deep learning, and computer vision are.

AI is the simulation of human intelligence using software and accompanying apps or machines. Machine learning, which is sometimes not included as a branch of AI since it is seen as its least robust form, automates analytical model building by identifying patterns in data and making decisions based on those perceived patterns.

Deep learning is a more robust machine learning technique that uses multi-layered hidden artificial neural networks in a "supervised" or "unsupervised" way, having as many as hundreds of network layers that either train on large amounts of data (supervised) or not (unsupervised) and then makes decisions on its own. Multi-layered neural networks are actually a system of mathematical algorithms that work together to come to a decision on what a pattern in a particular dataset is.

If deep learning is supervised, it uses labeled data as training inputs, that is, data that has been identified, collected together, and labeled according to particular attributes, such as visual data labeled as "car," "traffic sign," or "storefront." The output of "supervised" deep learning is the determination of whether or not the new data that is fed to it matches the training data in terms of particular attributes.

If the deep learning is unsupervised, it means that there is no labeled training dataset and the outcomes are unknown. The system learns on its own using its mathematical algorithms to see patterns. It is the unsupervised version of deep learning that has caught everyone's imagination. Both supervised and unsupervised deep learning are used for autonomous vehicles, but it is unsupervised learning that enables the vehicle to make a decision as to whether it should change lanes or suddenly stop. There is also a reinforced version of learning that mixes both supervised and unsupervised methods, but it is still the unsupervised "leg" that produces the outcomes that most mimic human decision-making.

Computer vision plays a very big part in both basic machine learning and deep learning. Corresponding to human vision, it serves as a mechanism by which both digital images and real-life images are “sensed,” with image attributes then sent to the learning mechanism. In this case, that is machine learning and deep learning; in the case of digital images, the use of computer vision is often called image processing. For instance, in the case of digital images, those that are in color are reduced to grayscale by the more advanced image processing algorithms, because this reduces the amount of data that the system has to deal with for identification. If color intensity is one of the attributes that needs to be categorized, then it is done by way of a digital tag. For real-life images, computer vision actually uses cameras that feed the visual data through its computer vision algorithms, which identify the elements of the images. This involves reducing everything to computer code that could then be fed through supervised machine learning and deep learning algorithms, first as training data if the data is more finely labeled by the computer vision algorithms, or, otherwise, fed through unsupervised deep learning algorithms.

In Spatial Computing, AI algorithms coupled with computer vision are used extensively. It is not hard to understand why this would be the case since Spatial Computing often deals with first identifying and then predicting the movement of three-dimensional objects. It can even work to identify and predict what you, a robot, or an autonomous vehicle will do next, or where you or it will look, within a Spatial Computing environment. This can be achieved because the AI algorithms, coupled with computer vision, produce a high level of situational “awareness” within the software that can then be used to make predictions about three-dimensional qualities and movements. New virtual beings can even walk around and interact with real and virtual items in your room, thanks to this technology.

Other examples of where digital situational “awareness” is necessary for Spatial Computing are when drones and autonomous vehicles are operated. For drones, the need for situational “awareness” goes beyond what is needed to have the drone move in space without colliding with other objects. Drones used for the delivery of packages, such as what Amazon is championing, as well as those used for warfare, have objectives that need to be fulfilled regarding the identification of objects and the predictions of correct trajectory that could be determined by AI algorithms and computer vision. Similar types of machines along these lines that benefit from using AI systems are food and parcel delivery robots that travel on city streets, as well as robots used in manufacturing and logistics that operate in irregular environments.

Another example of a machine that benefits from using AI systems is an autonomous vehicle. There are many similarities between the smaller delivery robot and an autonomous vehicle in terms of them having to seamlessly navigate city streets; however, the complexities associated with the rules of the road and the more sophisticated machinery and human payload warrant a much more robust deep learning system, one that incorporates computer vision that uses reinforced learning.

Outside of these physical manifestations using AI systems are those that integrate the virtual with the textual.

Machine learning really started to hit its stride when it became the system of choice for the user interface operation of bots, those text-based, mostly reactionary “assistants” that companies ranging from insurance to telecom, to retail to entertainment, and others are still using on their website direct messaging systems and for customer phone calls.

Virtual assistants, such as Siri, Amazon Alexa, and Google Assistant, are just vocalized bots that have been encapsulated within branded machinery, such as the iPhone, Amazon Echo, and Google Nest Hub. Where these become relevant to Spatial Computing is when these voice assistants are paired up with devices such as the Apple Vision Pro and AR glasses (to a lesser degree, they are currently useful spatially when used in collaboration with a smartphone and a mobile AR app). Voice navigation and command, paired with natural hand manipulation, will now serve as computing’s user interface.

Characters that embody a Generative AI bot virtually are called virtual beings. These virtual beings can be made three-dimensional, and thus spatial, and can be used in VR, as well as AR apps and experiences. Uses of these go beyond entertainment, into sales and marketing, training, customer relations, and other business uses—basically, anywhere where an expert is needed to provide relevant information.

In the further future, AI using unsupervised algorithmic systems could potentially create a culture and government for its own embodied bots, such as robots, which would ideally be in the service of humanity. These beings would be the ultimate in what Spatial Computing could bring. We choose to be optimistic with regard to this potentiality.

Back to the present. Since many tasks in Spatial Computing cannot be accomplished without the use of AI algorithms, often coupled with computer vision, AI’s impact on the business landscape cannot be overemphasized.

Evolution

The seven industry verticals that we will be covering in this book—1. Transportation; 2. **Technology, Media, and Telecommunications (TMT)**; 3. Manufacturing; 4. Retail; 5. Healthcare; 6. Finance; and 7. Education—are in the beginning stages of being transformed by AI-enhanced Spatial Computing. We will go into more detail on these transformations when we address each vertical in its individual chapter.

The evolution of personal computing, from a text-based interface to a three-dimensional one with Spatial Computing, fulfills our need for directness and ease that is produced via innovation. Spatial Computing's 3D interface jumpstarts a new higher curve for productivity.

The current technologies, such as Generative AI, that are significant to Spatial Computing will inevitably be enhanced and deepened with new technologies, all of which will take computing to new places that we cannot even imagine at this time.

In our next chapter, we discuss how Generative AI impacts Spatial Computing.

References

1. The 2024 WWDC can be viewed here: https://www.youtube.com/live/RXe0iIDNNek?si=KWFGe8w_JgYWGPkb

3

The Impact of Generative AI on Spatial Computing

Spatial Computing encompasses a fusion of physical and digital spaces where human interactions are facilitated by technology capable of understanding and leveraging the spatial relationships between objects. This sphere of technology intersects with various fields such as AR, VR, **Internet of Things (IoT)**, and robotics, among others. Generative AI, which is the subset of AI focused on creating new content—ranging from text to images to simulation environments—is poised to significantly impact Spatial Computing.

Enhancements in Immersive Experiences

Enhancements in immersive experiences through the application of Generative AI in AR and VR technologies represent a seismic shift in the simulation and training domains, as well as in entertainment and beyond. The realism and responsiveness that Generative AI introduces to virtual environments have multifaceted implications, particularly for professional training programs.

Hyper-Realistic Simulations

The advancements in Generative AI have ushered in an era of hyper-realistic simulations, particularly within VR training programs, that offer an immersive and effective means for professional development and education in various high-stakes fields.

The term “hyper-realistic” refers to a level of detail and accuracy that is as close to real life as possible. In the context of VR simulations, Generative AI is instrumental in achieving this level of realism.

It is capable of crafting high-resolution textures that mimic real-world surfaces down to the grain, producing intricate models that stand up to close inspection, and simulating physics interactions with high precision. This attention to detail is crucial in training programs where the nuances of the environment can significantly impact the trainee's learning experience.

For pilots, the implications are profound. Training in a VR environment that replicates every aspect of an aircraft's cockpit—from the resistance of a control yoke to the visual feedback from an instrument panel—can greatly enhance muscle memory and situational awareness. Such simulations can also replicate a wide range of scenarios that might be too dangerous, impractical, or rare to train for in real life, such as system failures, extreme weather conditions, or emergency landings. This can prepare pilots for the unexpected in a safe and controlled environment, which is invaluable for real-world application.

AI pilots have also been trained using simulated battles. In 2020, an AI pilot “won” a simulated dogfight against a real pilot¹.



Figure 3.1: VR pilot training program for the US Air Force and Navy at Eastern Kentucky University. This does not use Generative AI (yet). Photo credit: Eastern Kentucky University.

Surgeons stand to benefit dramatically from hyper-realistic simulations as well. VR training environments can provide an interactive platform where surgical procedures can be practiced on anatomically accurate virtual patients. These simulations can present a variety of patient cases, each with its own unique set of challenges and complications, offering a breadth of experience that would be impossible to gain quickly in the real world. Furthermore, the ability of Generative AI to simulate physiological responses—such as bleeding, changes in heart rate, or tissue reactions—to surgical interventions adds a layer of realism that can significantly improve the skills of a surgeon.

The realism in simulations is not just visual; it also extends to the physical laws that govern interactions within the simulations. When a trainee pilot executes a maneuver, the AI can generate realistic G-forces and aerodynamic feedback. In a surgical simulation, tools and tissues can exhibit accurate tactile feedback and responses. Such interactions are crucial for developing a deep, intuitive understanding of the subject matter, whether it's the aerodynamics of an aircraft or the texture and resistance of human tissue.

Generative AI is advancing to a point where it can introduce an emotional and psychological dimension to simulations. For instance, in a flight simulation, AI can generate scenarios with passengers and crew, complete with potential panic or emotional distress in emergency situations. This prepares pilots for the human element of flying, which can be just as unpredictable as the mechanical aspects.

Emotional Engagement and Behavioral Training

Training in a virtual environment often lacks the emotional and psychological pressures found in real-life situations. Generative AI can bridge this gap by simulating not only the visual and physical aspects of an environment but also the human elements. This includes generating virtual characters that exhibit complex, realistic behaviors. For example, training programs for law enforcement or military personnel can include AI-generated characters that display fear, aggression, or distress, training users to respond appropriately to human behavior.

Generative AI can also revolutionize how performance is assessed in training simulations. By analyzing a trainee's actions, the AI can provide tailored feedback, highlighting areas of strength and identifying areas needing improvement. This could be particularly beneficial in fields like aviation, where nuanced maneuvers are critical, and precision is non-negotiable. AI can track a pilot's eye movement, hand coordination, and reaction times, offering detailed insights that can be used to refine their skills.

The risks involved in training for high-risk occupations can be significant, with mistakes potentially leading to loss of life or significant property damage. VR and AR simulations, enhanced by Generative AI, offer a risk-free training environment. The cost of setting up physical training scenarios—like mock battlefields or full-scale airplane simulators—is also substantial. Generative AI can create these environments virtually, at a fraction of the cost, allowing for more frequent and accessible training sessions.

To sum it up, the integration of Generative AI with AR and VR is creating a new horizon in immersive experiences. It's not just about visual and auditory fidelity but also about creating comprehensive environments that engage all senses and emotions, adapting in real time to provide a deeply engaging and effective training experience. This technology offers the promise of safer, more efficient, and more thorough preparation for professionals in high-stakes fields. However, harnessing this power responsibly will require careful consideration of its psychological effects and the ethical implications of such realistic simulations.

Customization and Personalization

Customization and personalization through Generative AI within the domain of Spatial Computing are radically transforming user experiences by tailoring digital interactions to the individual preferences and behaviors of users. The implications of this technology are far-reaching and impact a myriad of sectors.

Interior Design and Architecture

The use of Generative AI in the field of interior design and architecture stands to fundamentally transform how living spaces and workspaces are designed, offering unprecedented levels of personalization.

Generative AI's ability to analyze and learn from individual preferences allows for a highly tailored design process. It can examine a user's previous selections of color palettes, furniture pieces, textures, and even responses to various architectural styles. With this information, the AI can construct a comprehensive profile of the user's taste and, subsequently, generate a personalized virtual representation of a space that aligns with their aesthetic sensibilities.

This technology enables an interactive and iterative design process where users can virtually "try on" different design elements in their own space using AR overlays. For instance, users can see how an Art Deco chandelier or a modernist sofa would fit into their existing living room, considering the play of natural light at different times of the day or the spatial flow between rooms. This mitigates the risk of design decisions that may look good in a catalog but not in the intended space.

Generative AI can also recommend spatial arrangements that optimize the functionality of an environment. This could mean suggesting furniture layouts that enhance movement within a room or employing smart storage solutions for space optimization. By understanding the user's lifestyle, Generative AI can make a living area not only aesthetically pleasing but also maximally efficient. It can account for the number of inhabitants, their activities, and the need for flexible spaces that can serve multiple purposes.

In architecture, AR using Generative AI can do more than just aesthetic simulation; it can predict the structural integrity of designs under various conditions. This allows architects to explore more innovative and daring designs, safe in the knowledge that they can virtually test the resilience of their structures. Examples of companies that are doing this include Trimble and Hololight.

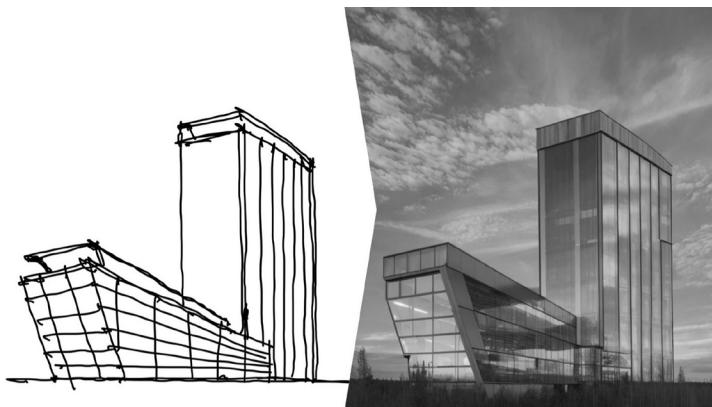


Figure 3.2: Using Generative AI, a sketch turns into an architecture render, which can also be viewed in 3D. Photo credit: Autodesk.

The integration of Generative AI in interior design extends to user experience by incorporating personal habits and routines into the design. For instance, a bedroom could be designed with automated window shades that adjust according to the user's sleep schedule, promoting better rest.

Looking ahead, the role of Generative AI in design is poised to grow. As the technology evolves, we could see AI becoming a collaborative partner in the design process, offering solutions that a human designer might not consider. It could potentially lead to a new design language that is at once profoundly personal, deeply functional, and ecologically conscious.

Generative AI represents a cutting-edge frontier in interior design and architecture, providing personalized and optimized design solutions that seamlessly merge aesthetics with practicality. As this technology continues to advance, it may redefine our relationship with our living spaces and workspaces, making them more responsive to our needs, preferences, and well-being.

Education and Training

The application of Generative AI to education and training is profoundly changing the landscape of learning, harnessing the potential to customize educational experiences in a manner that addresses the unique needs, preferences, and learning styles of each student.

Generative AI can design personalized learning journeys that adapt in real time to the progress and engagement of each student. It can analyze a student's interactions and performance to provide customized pathways through the curriculum. For example, if a student excels at algebra but struggles with geometry, the AI can adjust the coursework to provide additional geometry exercises while continuing to challenge them with more complex algebraic concepts.

Different students have different learning styles, and Generative AI can create multi-modal learning environments within VR and AR that cater to these preferences. Visual learners might be presented with rich graphical representations of concepts, such as 3D models of historical artifacts or animated illustrations of scientific processes. Auditory learners, on the other hand, could benefit from detailed narrations and discussions, while kinesthetic learners could engage with interactive elements that allow them to "learn by doing," performing virtual experiments or constructing models.

The interactivity and dynamism of content created by Generative AI can significantly enhance engagement and retention of knowledge. In VR and AR classrooms, historical events could come to life through reenactments or complex mathematical problems could be visualized and manipulated in three dimensions². This interactive content is not just passive but responds to the student's input, adjusting its level of complexity in response to their understanding.

Generative AI can provide continuous assessment and immediate feedback, an essential aspect of effective learning. By evaluating the student's performance in real time, the AI can identify areas of difficulty and automatically provide additional resources or alternative explanations. This immediate feedback loop enables students to learn from mistakes and misconceptions as they occur, which is a powerful pedagogical tool.

One of the most significant advantages of using Generative AI in education is the scalability of tailored education. Traditionally, personalized education required significant resources, often limiting it to one-on-one tutoring. With Generative AI, individualized learning experiences can be provided to a large number of students simultaneously, each receiving an education that adapts to their specific needs and pace.

Generative AI can also contribute to social and emotional learning by simulating social interactions and providing environments where students can practice communication, collaboration, and empathy. For instance, it could simulate a group project where students must work together to solve a problem, teaching them the value of teamwork and diverse perspectives.

As lifelong and life-wide learning becomes more important, Generative AI can support learners beyond traditional school settings, providing educational opportunities that fit into the learner's life schedule, adapting to changing career needs, and facilitating the acquisition of new skills throughout adulthood.

The challenges in implementing Generative AI in education include ensuring equal access to technology, protecting student data privacy, and preventing the entrenchment of biases in AI algorithms. Additionally, the role of teachers remains crucial; AI should be seen as a tool to augment, not replace, the human element in education.

The future of education with Generative AI promises a more adaptive, responsive, and engaging learning experience for students of all ages. As the technology continues to evolve, we can anticipate even more sophisticated and nuanced educational environments, with the potential to bridge gaps in understanding and cater to a diversity of learning needs like never before.

Generative AI presents an opportunity to reimagine education and training, offering personalized, engaging, and effective learning experiences. By catering to individual learning styles, providing interactive content, and adapting to each student's pace, Generative AI is not just enhancing education—it is transforming it to be more inclusive, adaptable, and aligned with the diverse needs of learners in a rapidly changing world.

Gaming

The gaming industry stands on the brink of a transformative era with the integration of Generative AI and Spatial Computing, which together can create deeply immersive and reactive gaming experiences that respond in real time to a player's decisions and play style.

Role-playing games (RPGs) can be redefined by Generative AI, as it allows the narrative to evolve dynamically, offering a storytelling experience that is as unique as the player. Instead of choosing from a set of pre-determined pathways, a player's actions, dialogues, and in-game decisions could organically shape the storyline.

A character's alignment, whether lawful or chaotic, brave or cautious, could influence the unfolding events and available quests in nuanced ways that reflect the player's journey.



Figure 3.3: A non-player character (NPC) that's part of the Angry Birds AR NPC pack. With Generative AI, an NPC can be generated and featured automatically within a game. Photo credit: Resolution Games.

Beyond narratives, Spatial Computing can change gameplay mechanics. In stealth-based games, if a player is adept at sneaky maneuvers, the game might introduce more complex security systems or AI opponents with heightened senses to provide a challenge appropriate to their skill level. Conversely, for those who prefer a more direct approach, the game could spawn tougher enemies suited for melee or ranged combat, demanding the player to refine their strategy and skills.

Generative AI can modify the game environment in real time, making it reactive to the player's strategies and preferences. If a player engages more with the environment by using objects as cover, the AI could generate more such interactive elements, altering the layout of future levels to match the player's tactical engagement. The virtual environment could also evolve based on in-game events. For example, a city might show signs of economic growth or decline as a direct result of the player's actions, affecting gameplay such as available resources or NPC behavior.

Games can use Generative AI to tailor challenges and rewards to the player's performance. If the game detects that a player is skilled in certain areas, it might offer them harder challenges with correspondingly greater rewards, thus maintaining engagement by ensuring that the game is neither too easy nor frustratingly difficult.

Generative AI can be used to create content such as quests, characters, and even whole worlds. This content is not static but evolves based on the collective actions of the player base. For instance, if players show a preference for certain types of quests or environments, the AI can generate more content along those lines, ensuring a fresh and relevant gaming experience.



Figure 3.4: Niantic's Peridot, an interactive AR game with AR pets. This game does not use Generative AI but imagine a game like this that does. Photo credit: Niantic.

Advanced Spatial Computing could allow for emotional and contextual interaction within games. Using voice recognition and analysis, a game could detect the player's emotional state and adjust the difficulty or narrative accordingly. This could create a more emotionally engaging and empathetic gaming experience, with NPCs responding differently depending on the player's tone or choice of words.

As Generative AI and Spatial Computing continue to mature, we can expect games that are not only reactive but also proactive, with AI that can anticipate player needs and preferences, potentially transforming the gaming industry with experiences that are incredibly personalized and engaging.

The intersection of gaming, Generative AI, and Spatial Computing promises a future where games are not just played but experienced in a way that is deeply personal and continuously evolving. This not only maximizes player engagement but also opens the door to new forms of storytelling and gameplay that have yet to be imagined, setting the stage for a new era in the gaming industry.

Beyond Gaming

Beyond gaming, the combination of Spatial Computing and Generative AI is especially fruitful when it comes to the Apple Vision Pro and its passthrough AR. With Generative AI, anyone can create visuals, including short-form videos, and look at them in the Apple Vision Pro, not only in 2D but in 3D. Yes, this could also be done on the Meta Quest 3, but the ease of use of the Apple Vision Pro is a big plus. This is not only for regular, everyday consumers but also for creators. So, both consumers and creators can truly benefit from the combination of Spatial Computing and Generative AI. This encompasses the personal arena of family and friends, and both amateur and professional creators. Hollywood is already on notice regarding Generative AI. The combination of Spatial Computing and Generative AI heralds a new era for Hollywood, promising to revolutionize both the creative and production processes involved in filmmaking.

Spatial Computing allows for the creation of virtual sets and environments, which directors and cinematographers can explore and manipulate in real time. Generative AI can fill these spaces with procedurally generated props, landscapes, or even crowds, speeding up the pre-visualization process. This enables filmmakers to experiment with different scenarios and visual effects before physical production begins, reducing costs and enhancing creative flexibility.

Generative AI can be used to create more realistic and complex **visual effects (VFX)**. For instance, it can simulate realistic environmental conditions, such as weather patterns or the movement of water, that would be costly or even dangerous to reproduce physically. AI can also generate lifelike textures and skins for CGI characters, making them more convincing to the audience.

Spatial Computing enhances motion capture technology, allowing actors' performances to be captured in more detail and translated into digital models with greater accuracy. Generative AI can then take this data to create nuanced animations, particularly for non-human characters, which can adapt dynamically to the scenes they are in, leading to more natural and responsive interactions within the film's narrative.

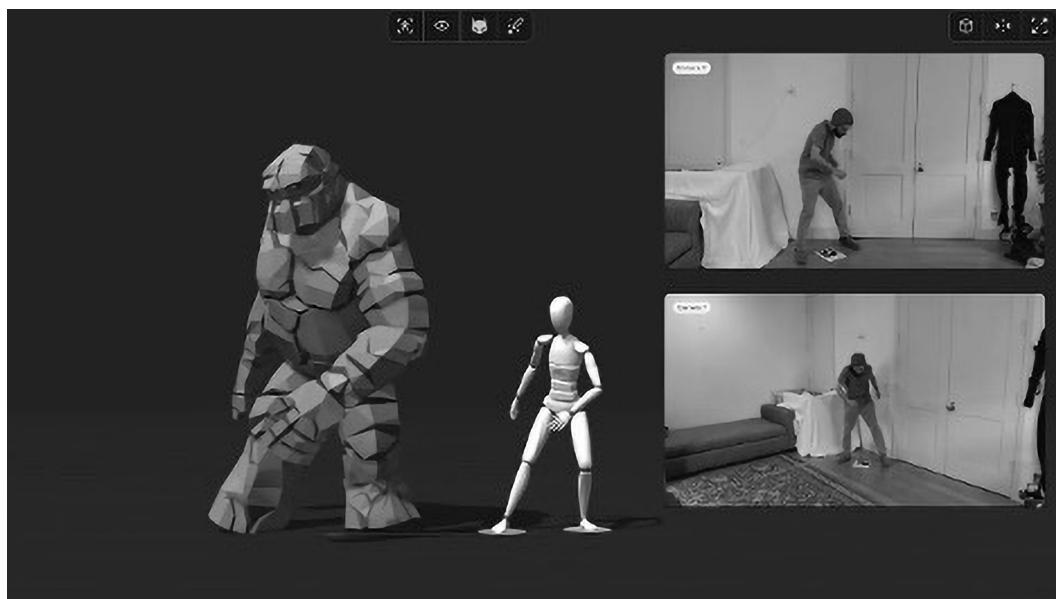


Figure 3.5: AI performance capture that can use Generative AI to generate a non-human character or anything else that you want. Photo credit: Rokoko Vision.

In interactive or immersive media, Spatial Computing and Generative AI enable the creation of films that respond to the viewer's gaze, gestures, or choices, providing a unique story experience for each viewer. This could lead to new forms of narrative that are less linear and more personalized, akin to a blend of film and video gaming.

Generative AI can assist in the editing room by automatically compiling raw footage based on parameters set by the editors. It can analyze takes for quality, continuity, and performance, suggesting the best options to the human editor and thereby streamlining the post-production process.

Generative AI can be used in costume design and makeup, offering a digital catalog of styles and effects that can be virtually tested on actors before physical production. This can extend to the creation of virtual wardrobes and makeup effects that are consistent with the world-building of the film.

Spatial Computing and Generative AI can be used for marketing films through personalized trailers and AR experiences that engage potential viewers by providing content tailored to their interests and past behavior. This technology could be used to analyze social media trends and viewer data to create targeted marketing campaigns that have a higher chance of resonating with audiences.

The fusion of Spatial Computing and Generative AI holds tremendous potential for Hollywood, offering innovative tools for storytelling, production, and viewer engagement. This synergy is poised to open up new creative avenues, redefine traditional filmmaking paradigms, and offer more immersive and personalized experiences for audiences worldwide. As these technologies continue to mature, they will undoubtedly shape the future of entertainment.

Retail and Shopping

Spatial Computing and Generative AI could also transform the retail experience. AR-enabled shopping apps can recommend products in real time as the customer moves through a store. These recommendations are not random; the AI generates them based on the customer's past shopping history, the sizes they prefer, the colors they wear most often, and even current trends they have shown interest in. In a virtual dressing room, the AI can superimpose clothing on the user's image, reflecting their style and fitting preferences, and even suggesting accessories to complement their selections.

One of the key advantages of Generative AI in Spatial Computing is the ability to provide personalization at scale. This means that instead of generic experiences designed for the average user, every individual can have a unique interaction that feels specially crafted for them. This has significant implications for customer satisfaction and can lead to greater brand loyalty in competitive markets.

Customization and personalization through Generative AI are not simply about enhancing user enjoyment or convenience; they represent a paradigm shift in how individual needs and preferences are understood and catered to within digital and spatial environments. As this technology advances, the scope of personalized experiences will continue to expand, potentially reshaping industries and redefining expectations for individualized interaction with technology. However, as with any technological innovation, it is crucial to proceed with a conscientious approach that prioritizes user agency and ethical data use.

Personalized Work Environments

In the world of work, especially with the trend toward remote and flexible work arrangements, the role of Generative AI and Spatial Computing in personalizing work environments is becoming critically important. These technologies can enhance not only the physical aspects of a workspace but also the social and interactive dynamics of working collaboratively from a distance.

Generative AI can analyze an individual's work patterns and environmental preferences to optimize the physical workspace. By employing Spatial Computing, a virtual office can be rendered where lighting adjusts in real time to reduce eye strain or to mimic natural light cycles, potentially boosting mood and concentration. Background noise can be managed to ensure an auditory environment that balances the need for quiet focus with the preference for ambient noise that can enhance cognitive performance for some individuals.

Virtual workspaces can be made adaptive to the ergonomic needs of the user. Generative AI could suggest desk and chair adjustments, monitor heights, and even recommend breaks for stretching based on the user's posture and time spent at the desk. Such a personalized approach to ergonomics could help reduce work-related musculoskeletal disorders and contribute to longer-term health and well-being.

For remote workers, VR and AR can offer a simulated office environment where interactions with colleagues are more natural and engaging. Generative AI can create avatars that not only resemble the user's physical appearance but also mimic their gestures and communication style, leading to more nuanced and effective virtual meetings. These avatars could be programmed to express nonverbal cues that reflect the user's real-life reactions, helping to bridge the gap between virtual and in-person communication.

Generative AI can play a pivotal role in task management by organizing and prioritizing workloads based on the user's work habits and peak productivity times. It could dynamically schedule tasks, reminding users of deadlines and suggesting the optimal sequence of activities to maximize productivity without leading to burnout.

A personalized virtual work environment could also support emotional well-being and mental health. AI could detect signs of stress or fatigue through user interaction patterns and suggest appropriate interventions, such as a short meditation session or a virtual walk in a peaceful environment.

For those working from home, Spatial Computing could integrate with smart home systems, allowing the work environment to interact with other home devices. For instance, the heating could turn on shortly before the workday begins, or smart blinds could adjust to reduce glare on the computer screen without the user having to take a break from their work.

As with all applications of AI that involve personal data, security and privacy are paramount. Systems must be designed to protect sensitive information, and users must have control over what data is collected and how it is used.

As Generative AI and Spatial Computing technologies continue to advance, the potential for creating highly personalized and responsive virtual work environments grows. These technologies could lead to workspaces that adapt not just to functional needs but also to the evolving personal development and well-being goals of each individual, fostering a more satisfying and productive work life.

The personalization of work environments using Generative AI and Spatial Computing represents a significant step forward in creating workspaces that are not only about work efficiency but also about the overall quality of work life. By catering to individual needs and preferences, companies can harness the benefits of increased employee satisfaction and productivity, paving the way for a future where work is more adaptable, humane, and responsive to the human elements of the workforce.

Future Directions for Customization and Personalization

As Generative AI technology continues to mature, we can expect even more sophisticated levels of customization and personalization. This could lead to the development of smart environments that not only adapt to individual users but also anticipate their needs and preferences, possibly even before the users are aware of them themselves.

The interplay of customization, personalization, Generative AI, and Spatial Computing is forging a new frontier in how we interact with technology. It presents opportunities for tailored experiences that heighten engagement, improve functionality, and increase satisfaction across various sectors. The full realization of these technologies' potential will depend on the continuous refinement of AI algorithms and a steadfast commitment to addressing the ethical considerations that accompany such personalized experiences.

Dynamic Content Generation in More Detail

Dynamic content generation via Generative AI has the potential to revolutionize Spatial Computing by providing an environment that is not just interactive, but also adaptive and responsive to the context and behavior of users. This has far-reaching implications for various industries and user experiences.

The incorporation of Generative AI into VR gaming is paving the way for a new generation of gaming experiences, characterized by procedurally generated environments that offer a level of variety and customization previously unattainable.

Generative AI enables the creation of infinite landscapes, with each terrain generated algorithmically rather than being manually crafted. Players can explore worlds that stretch far beyond the confines of pre-designed levels, discovering unique biomes, terrains, and ecosystems. This endless exploration captures the essence of adventure, as players know that beyond every horizon lies a landscape that even the game's developers have not seen.

Procedural generation also allows for environments to be tailored to a player's skill level and playing style. If a player thrives in puzzle-solving, the AI might generate intricate labyrinths or environmental puzzles that require thoughtful navigation. For those who prefer combat, the terrain could organically form arenas suited for battles, or for the strategist, landscapes that lend themselves to defense and fortification.

Challenges and obstacles within these environments can be generated and positioned based on the player's progress. This dynamic content generation ensures that gameplay remains engaging over time. Generative AI can create these obstacles to provide just the right level of difficulty, adapting to the player's learning curve to keep the game challenging but not discouraging.

Generative AI can be programmed to consider the player's narrative choices, adjusting the environment to reflect the evolving story. For example, a player's actions in a game could lead to the restoration or decay of a virtual city, with the environment dynamically changing to reflect the state of the world according to the player's decisions and actions.

As a player engages with the game world, the AI can use their interactions to evolve gameplay mechanics and environmental elements over time. A player who frequently engages in exploration might find new paths and hidden areas emerging, while one who builds and crafts might see their structures influence the development of the world around them.

In games that simulate societies or ecosystems, Generative AI can produce complex, functioning systems that react to the player's influence. The behavior of AI-controlled characters and wildlife can adapt to environmental changes, leading to emergent gameplay where the player's actions ripple through the game world in unpredictable ways.

Procedural generation not only increases replayability but also fosters community engagement as players share unique experiences and discoveries that are exclusive to their playthrough. It becomes a shared yet personal journey, as each player can contribute stories of their unique adventures.

The unpredictability of procedurally generated environments can be both a boon and a challenge. Game designers must ensure that generative algorithms produce coherent and navigable spaces while preserving the sense of wonder and discovery. Looking to the future, we may see AI that not only generates the environment but also learns from the collective experiences of the player base to refine and enhance the generation process.

Generative AI represents a significant step forward for VR gaming, moving away from the limitations of pre-designed environments to offer players a truly unique and personal journey. With procedural generation, the landscapes, challenges, and stories of a game become as dynamic and limitless as the players themselves, heralding a new era of personalized and enduring gaming experiences.

Adaptive AR Applications

Adaptive AR applications enhanced by Generative AI represent a transformative development in how information is presented and consumed in the physical world. By combining real-time data analysis with spatial awareness, these applications can deliver personalized content that significantly alters the user experience in various contexts.

In advertising, adaptive AR applications could mark a paradigm shift from passive reception to active engagement. Generative AI can tailor advertisements in real time based on a user's browsing history, purchase records, or even mood, which can be inferred from mobile device data or wearable technology. Digital billboards might change as a person walks by, showing an ad for a product they recently searched for online or a deal at a nearby restaurant they frequent. This personalization could dramatically increase the effectiveness of advertising by presenting relevant offers and information, thereby increasing conversion rates.



Figure 3.6: A Generative AI advertisement put out by Coca-Cola. This advertisement could then be made into an AR image that could be viewed on a Spatial Computing headset such as the Apple Vision Pro. Photo credit: Coca-Cola.

In retail, AR can enrich the shopping experience by providing customers with personalized recommendations as they browse through stores. Generative AI could generate AR overlays on a shopper's smartphone or smart glasses, highlighting products similar to their past purchases or items that complement something they have already picked up. It could also offer virtual try-ons, display additional product information and reviews, or even show how a product could be used or worn.

Museums and exhibitions could use adaptive AR to create interactive learning experiences that respond to the interests of each visitor. As a visitor engages with an exhibit, Generative AI could present additional information on topics of interest, suggest pathways through the museum, or provide interactive experiences that go deeper into subjects the visitor has spent more time exploring.

Adaptive AR applications could revolutionize navigation by offering route guidance that takes into account not just the destination but also the user's preferences, such as avoiding crowded areas or including scenic paths. Additionally, AR can overlay historical facts, architectural information, or even reviews of restaurants and shops as a user views them, creating an enriched layer of information atop the physical world.

In smart cities, adaptive AR can be leveraged to enhance citizen interactions with public services. Generative AI could provide residents with information about city events, public transportation updates, or local news relevant to their location and preferences. This real-time information can be spatially anchored to physical locations, allowing residents to visually connect data with their urban environment.

As Generative AI and AR technology continue to advance, we can anticipate even more sophisticated applications. For example, AR could provide contextual information to workers in complex tasks, adapting to their skill level and learning pace. In social settings, AR might provide conversation prompts or background information to help individuals connect more meaningfully with those around them.

Adaptive AR applications driven by Generative AI can create deeply personalized and engaging experiences that merge the digital and physical worlds. This technology stands to revolutionize a range of industries, from advertising to education to urban planning, offering tailored content that enhances interaction and utility. As this field progresses, maintaining a balance between personalization and privacy will be paramount to its acceptance and success.

Dynamic content generation can also extend to user interfaces in Spatial Computing, where elements of the interface change in response to the user's behavior, preferences, or context. For example, an AR application could change its menu layout to prioritize the most commonly used functions or adapt its presentation of information based on the user's interaction history, thus enhancing usability and efficiency.

Real-Time Data Visualization

The integration of Generative AI with Spatial Computing in the context of data analysis and visualization opens up dynamic new ways for stakeholders to interact with and understand complex datasets in real time.

Generative AI can construct immersive 3D models of environments such as cities, where each element is mapped to live data streams. This creates a living, breathing virtual replica of the city that officials, planners, and residents can explore in VR.

Such models could represent traffic flows, showing congestion in real time, or the movement of public transportation, allowing users to spot delays or inefficiencies immediately.

Live weather data can be visualized in this virtual environment, showing rain clouds forming, temperatures changing, or even pollution levels rising in different areas. This kind of visualization goes beyond traditional maps or charts, offering an intuitive understanding of environmental conditions and their potential impacts on the city.

Real-time population data can also be incorporated, perhaps showing the density of crowds in different districts, allowing city planners to allocate resources or manage events more effectively. Utilities such as water and electricity distribution could be monitored, with AI highlighting areas of heavy usage or potential faults in the system before they cause outages.

Generative AI can enable interactive analytics within these environments. Users can pose questions or run simulations to see the potential effects of their decisions. For example, if city officials are considering a new public park, they could visualize how it might affect local traffic, pedestrian flows, or even the microclimate of the area.

The dynamic visualization of data can also offer insights into economic activity and social behavior. Businesses could use this technology to identify consumer patterns or to pinpoint optimal locations for services. Social services could use it to track community health trends or to optimize the distribution of aid.

By making data visualization more accessible and interactive, such VR and AR models could be used to engage the public, educate students, or train new employees. The level of detail and interactivity would make the data more relatable and understandable, potentially increasing civic engagement and awareness.

Generative AI allows these visualizations to be customizable and scalable. Users could zoom in on particular data points of interest or expand their view to understand regional trends. The system could even adapt to user interactions, highlighting areas that many users are interested in or providing additional layers of data in response to common queries.

The challenge with real-time data visualization lies in ensuring the accuracy and security of the data being represented. Misrepresentation could lead to incorrect decisions with real-world consequences. Additionally, the complexity of creating such systems is significant, requiring advanced AI capable of processing and visualizing vast datasets quickly and reliably, which also carries with it great financial expense.

Generative AI's role in dynamic content generation within Spatial Computing environments promises to create experiences that are not just personalized, but also inherently capable of growth and change in sync with the user. The potential of such technology is vast, with applications that can enhance entertainment, personalize learning, and transform data interaction. As the technology matures, it will be crucial to navigate the challenges of privacy, ethics, and the balance between personalization and exposure to diverse content, to fully realize the benefits of dynamic content generation in Spatial Computing.

Accessibility

The application of Generative AI within Spatial Computing holds transformative potential for enhancing accessibility for individuals with disabilities. By leveraging the ability to interpret and respond to three-dimensional spaces, Generative AI can provide tools and solutions that were previously unimaginable, thus leveling the playing field in many aspects of life for those with disabilities.

For the deaf and hard-of-hearing community, real-time virtual sign language interpreters created through Generative AI could be a game-changer. Such AI systems could interpret spoken language and translate it into sign language instantly, with avatars or holograms performing the signs. This would facilitate communication in a variety of settings, from lectures and public speeches to everyday conversations. By incorporating this into AR glasses, for instance, a user could have their own personal interpreter accompanying them wherever they go, making communication seamless and more inclusive.

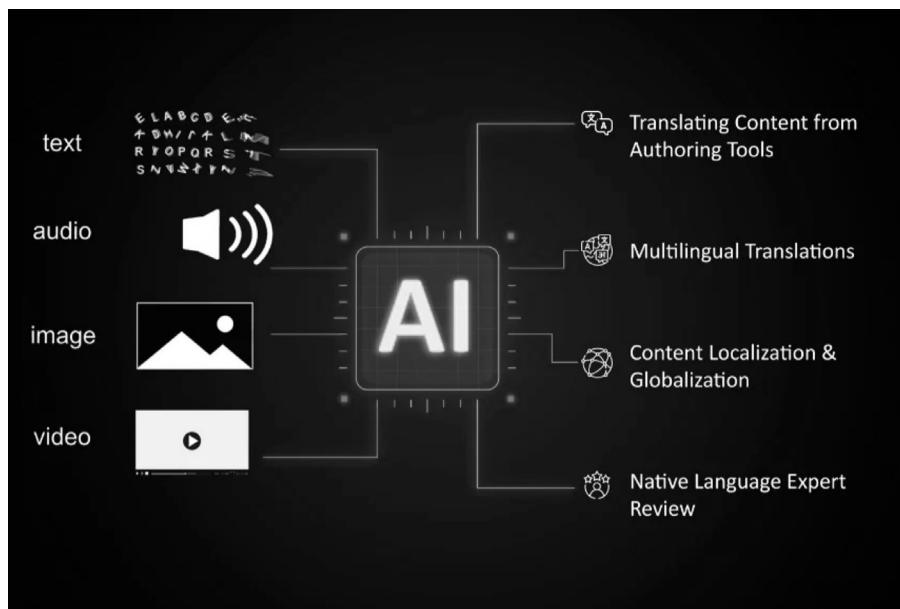


Figure 3.7: Real-time sign language interpretation uses computer vision to identify movements and Generative AI to interpret words and then use sound to sound out the words. Photo credit: Harbinger Group.

Navigational aids for visually impaired individuals are another domain where Spatial Computing and Generative AI can significantly enhance accessibility. By understanding the layout of an environment, AI can provide auditory cues and haptic feedback through AR devices to help visually impaired users navigate spaces safely and independently. For example, a spatially aware AR application could inform a user of the proximity and nature of objects in their path, suggest clear routes, warn of upcoming stairs or uneven surfaces, and even describe their surroundings to create a mental map of the space.

Generative AI can also create customizable interfaces in VR or AR that adapt to the physical abilities of users with motor impairments. Such interfaces could reconfigure themselves to be operable with minimal movement or to respond to alternative input methods, such as eye tracking or voice commands. This customizability would allow individuals with a wide range of physical disabilities to interact with digital content and control their environment in a way that is tailored to their capabilities.

In educational settings, Spatial Computing aided by Generative AI could dynamically adapt content to be accessible to students with various disabilities. Text could be automatically converted into braille for tactile display, or lessons could be transformed into formats that are more suited to students with learning disabilities, like dyslexia, by using symbols, interactive 3D models, and auditory storytelling.

Generative AI can facilitate the development of AR applications that assist individuals on the autism spectrum with social interactions. When analyzing facial expressions and body language in real time, these applications could provide cues or advice to the user, helping them interpret social signals that they might otherwise find challenging to understand. This can support them in social learning and reduce anxiety in social settings.

Spatial Computing, empowered by Generative AI, can greatly expand accessibility for individuals with disabilities, enabling more autonomy and enhancing the quality of life. From navigation aids to real-time translation of sign language, the potential applications are vast and varied. As with all technological advancements, though, the deployment of such technologies must be guided by ethical considerations, user-centric design principles, and a commitment to protecting personal data. The future of accessible technology through Generative AI and Spatial Computing looks promising, with the potential to create a more inclusive world for all.

Data Synthesis and Simulation

Data synthesis and simulation through Generative AI in Spatial Computing present significant opportunities for enhancing predictive modeling and decision-making in various fields, particularly in urban planning and environmental studies.

Urban planning requires the consideration of countless variables—from traffic flow and public transportation to zoning and land use. Generative AI can create multi-layered simulations of urban environments that synthesize data from various sources. Planners can input population growth estimates, proposed building plans, traffic data, and other relevant statistics into a Generative AI model, which can then simulate how these factors will interact over time. This process results in a detailed visualization of potential urban development paths, allowing planners to observe the long-term impacts of their decisions. By doing so, they can identify the best strategies for sustainable growth, optimize traffic management, and plan infrastructure effectively.

In environmental science, Generative AI can simulate the effects of climate change, urbanization, and other factors on natural ecosystems. For example, researchers could use Spatial Computing to visualize the progression of deforestation, its impact on local biodiversity, and the potential effects of various conservation strategies. These simulations can extend to predicting the spread of wildfires, the impact of rising sea levels on coastal communities, or the potential pathways of invasive species.

Disaster response strategies can be greatly improved with Generative AI simulations. Spatial Computing can model scenarios such as floods, earthquakes, and storms within virtual environments, incorporating geographical data, historical disaster patterns, and current urban layouts. Emergency services can use these simulations to plan evacuation routes, determine the locations where resources should be deployed, and conduct disaster response drills. By anticipating the various ways a disaster could unfold, responders can be better prepared to save lives and property when a real disaster strikes.

Generative AI can also be used to simulate economic scenarios within a spatial context. It can model how changes in policy, such as tax incentives or zoning laws, could affect local economies. By integrating data on consumer behavior, employment statistics, and business activity, these simulations can help economists and policymakers predict the outcomes of their decisions with greater accuracy, potentially leading to more prosperous and resilient communities.

Spatial Computing and Generative AI are highly applicable in health and epidemiology, where they can simulate the spread of diseases across different geographical landscapes. Taking into account factors such as population density, mobility patterns, and public health interventions, these simulations can provide valuable insights into how infectious diseases could spread and how different containment strategies might alter the trajectory of an outbreak.

Generative AI has the potential to vastly enhance the field of data analysis and predictive modeling within Spatial Computing environments. By synthesizing large datasets and simulating complex systems, it enables more informed decision-making across a spectrum of disciplines. These advanced simulations can lead to more resilient urban planning, better environmental management, improved disaster response, deeper economic insights, and more effective public health strategies. As Generative AI continues to evolve, its role in data synthesis and simulation will become increasingly integral to our ability to plan for and respond to the complex challenges of our world.

The data synthesis and simulation capabilities of Generative AI and Spatial Computing come with challenges that must be addressed to ensure their effective and ethical use. The accuracy of simulations depends heavily on the quality and breadth of the data fed into them. Bias in data can lead to inaccurate or misleading outcomes, so careful validation and constant refinement of AI models are necessary. There's also the risk of over-reliance on simulations for decision-making; they should complement, not replace, human judgment and empirical research.

Challenges and Ethical Considerations

As the fusion of Generative AI and Spatial Computing continues to advance, it presents a host of challenges and considerations that must be navigated to harness its full potential while maintaining ethical standards and real-world applicability.

Data Quality and Bias

Data quality and bias are indeed pivotal in shaping the reliability and fairness of applications that leverage Generative AI and Spatial Computing. These two factors directly influence the performance and outcomes of AI systems, with a significant impact on real-world decisions and perceptions.

Data quality pertains to the suitability of data for making decisions. High-quality data must be accurate, complete, timely, consistent, and collected using robust methodologies. For Generative AI systems, which often learn and make inferences from large datasets, the adage “garbage in, garbage out” is particularly apt.

The use of poor-quality data can lead to inaccurate models that, in turn, produce unreliable or misleading results. For instance, if a Spatial Computing system is tasked with optimizing a public transportation network but is fed incomplete data, it might fail to suggest viable improvements or might propose solutions that could exacerbate existing issues.

Bias in AI can arise at any stage of the data life cycle—from collection to processing and interpretation. It can take many forms, such as:

- **Sampling bias:** This occurs when the data collected is not representative of the entire population or scenario. For example, a Generative AI system designed to simulate pedestrian traffic patterns must have data that accounts for all demographics to avoid urban designs that cater only to certain groups.
- **Measurement bias:** This takes place when the data collection instruments or procedures influence the data. An AI system trained to recognize speech patterns must be exposed to a variety of accents and dialects to ensure it does not favor one group over another.
- **Confirmation bias:** This happens when the data is interpreted or selected in a way that confirms pre-existing beliefs or hypotheses. For Spatial Computing applications used in predictive policing, for instance, this can lead to over-policing in certain areas, perpetuating cycles of discrimination.

The consequences of biased AI systems can be profound. In the context of Spatial Computing, biased data can lead to the creation of environments that do not serve the needs of the entire population. In scenarios like disaster response, this could result in inadequate services to those most in need, or in the context of resource allocation, it could perpetuate socioeconomic disparities.

Mitigating the risks associated with data quality and bias requires a multi-pronged approach:

- **Diverse data collection:** Ensuring that datasets are diverse and representative of different populations and environments. This might involve gathering data from a variety of sources and contexts.
- **Regular data auditing:** Implementing continuous audits of data and models to identify and correct issues of quality and bias. This might be carried out by independent third parties to ensure objectivity.
- **Algorithmic transparency:** Providing clear explanations of how AI systems make decisions. This can help stakeholders understand and trust AI-driven processes and outcomes.
- **Ethical AI frameworks:** Establishing ethical guidelines that govern AI development and deployment, including considerations for privacy, consent, and the avoidance of harm.

- **Interdisciplinary teams:** Bringing together experts from different fields—including those who understand the technical aspects of AI and those who have insights into the ethical and social implications—to guide the development of these systems.

As Generative AI and Spatial Computing continue to evolve, so will our understanding of data quality and bias. Future research and development efforts must focus on creating more sophisticated methods for assessing data quality and detecting bias, as well as developing AI systems that can adapt to and compensate for these issues. The goal is to ensure that AI-driven applications are both technically sound and socially responsible, fostering trust and advancing equity in all areas where they are applied.

Ethical Considerations

The ethical implications of Generative AI in Spatial Computing extend beyond data integrity. There are privacy concerns, especially as AI systems become more adept at synthesizing information that may inadvertently reveal personal details about individuals or sensitive information about organizations. There's also the moral responsibility of deploying such simulations in a manner that does not inadvertently harm certain populations or the environment.

The integration of Generative AI with Spatial Computing technologies has the potential to significantly impact individual privacy. AI systems that collect and process large volumes of data to create personalized experiences could inadvertently expose sensitive personal information. For example, an AR application might reveal a person's location, habits, or even health information through the data it collects. Ensuring that data is used responsibly and with the explicit consent of users is paramount. Furthermore, there should be strict controls over how data is stored, who can access it, and how long it is retained.

Generative AI's ability to synthesize information could lead to the inadvertent disclosure of proprietary or confidential information about organizations. This could occur through the accidental inclusion of sensitive data in training datasets or the AI's ability to infer and reconstruct protected information. Robust data governance and rigorous de-identification protocols are necessary to prevent such disclosures.

The deployment of AI-driven simulations could have unintended consequences for certain populations. For instance, an urban planning tool might inadvertently prioritize the needs of one demographic over another, leading to social inequities. There's also the potential for environmental harm if AI simulations lead to decisions that negatively impact ecosystems or contribute to unsustainable practices. AI developers and users must consider the broader societal and environmental ramifications of their work, striving for inclusivity and sustainability.

Generative AI may encode and perpetuate existing biases, leading to unfair outcomes. For example, a Generative AI system used in job recruitment might favor candidates from a particular demographic background if the training data reflects historical biases. Ensuring algorithmic fairness involves actively identifying and correcting biases in both datasets and AI models.

There is a need for transparency in how Generative AI systems make decisions and adapt over time. Users should be able to understand the basis on which information is being presented to them, especially when it influences decision-making in critical areas such as healthcare, law enforcement, or financial services. Accountability structures must be in place to address any harm caused by AI systems.

The risk of over-reliance on AI-generated simulations is a significant concern. While simulations can provide valuable insights, they are, by nature, simplifications of reality. Complex human behaviors, random environmental factors, and unpredictable global events are difficult to model accurately. Over-reliance on simulations could lead to a lack of preparedness for unexpected outcomes or a false sense of security in the face of complex issues. It is vital to use simulations as one of several tools in decision-making, not the sole basis.

The current state of technology imposes limitations on what Generative AI and Spatial Computing can achieve. Processing power, data storage capacity, and the ability to integrate different data sources are all factors that can limit the scope and accuracy of simulations. Advancements in quantum computing and more sophisticated AI models may overcome some of these limitations in the future. Advancements in quantum computing can address current limitations in Generative AI and Spatial Computing by offering significantly enhanced processing power and data storage capabilities. Quantum computers can perform complex calculations much faster than classical systems, allowing for more sophisticated simulations and real-time data processing. Additionally, quantum algorithms can improve data integration and AI model training, leading to more accurate and efficient Generative AI outputs. These advancements will help overcome existing constraints, enabling more powerful and comprehensive applications.

Integrating human expertise with Generative AI is essential. Professionals from various domains—be it urban planning, environmental science, or public health—can provide the contextual understanding and interpretive skills that AI currently lacks. This collaboration can lead to more robust and realistic simulations.

Educating Stakeholders

For Generative AI and Spatial Computing to be effectively implemented, stakeholders, including policymakers, industry leaders, and the general public, need to be educated about their capabilities and limitations. Misunderstandings about what AI can and cannot do can lead to misuse or distrust of the technology.

Developers and deployers of Generative AI systems have a moral responsibility to consider the ethical dimensions of their work. This includes not only the direct outcomes of AI applications but also their long-term implications for society. Ethical principles such as beneficence, nonmaleficence, autonomy, and justice should guide the development and application of these technologies.

Engaging in public discourse about the ethical use of Generative AI in Spatial Computing is crucial. This includes not only experts in the field but also the public, policymakers, ethicists, and other stakeholders. Regulatory frameworks may need to be developed or updated to ensure that the use of these technologies aligns with societal values and ethical norms.

The ethical considerations of Generative AI in Spatial Computing require proactive engagement and ongoing assessment. As the capabilities of these technologies grow, so too does the need for a principled approach that respects individual rights, promotes fairness, ensures accountability, and safeguards public trust. Addressing these ethical challenges is critical for the responsible advancement and integration of AI in society.

Future Directions

Looking forward, the field is likely to evolve in ways that address these challenges. This could include developing new forms of data encryption to protect privacy, creating industry-wide standards for AI simulations, and investing in educational programs that build AI literacy. Additionally, the continued refinement of AI algorithms, coupled with an increasing emphasis on interdisciplinary collaboration, promises to enhance the accuracy and utility of simulations in Spatial Computing.

Generative AI and Spatial Computing are at the forefront of a technological revolution with the power to transform industries and improve lives. However, their successful implementation hinges on addressing current challenges through rigorous validation, ethical oversight, and a balance between technological reliance and human judgment. As these technologies progress, they present a promising avenue for innovation, provided they are steered with caution and consciousness of their broader impact on society and the environment.

The integration of Generative AI with Spatial Computing is transforming how we interact with digital information and virtual environments. The ability to create personalized, dynamic, and realistic experiences holds tremendous potential for a variety of sectors, including entertainment, education, healthcare, and urban development. However, to ensure these advancements lead to broadly beneficial outcomes, it is essential to address the accompanying ethical and practical challenges head-on, establishing robust frameworks for responsible use.

In the future, as Generative AI continues to evolve and as Spatial Computing becomes even more integrated into daily life, we can expect a continuous redefinition of human-computer interaction paradigms. The synergy between these technologies will likely spur innovation, yielding new tools and applications that we can scarcely imagine today. It is an exciting time for developers, consumers, and society at large, as we witness and participate in the unfolding of this digital revolution.

In our next part, *Part II*, we provide you with a roadmap for the disruptions caused by Spatial Computing that will happen in seven industries, starting with the transportation industry.

References

1. A dogfight between AI and a human pilot: <https://news.sky.com/story/first-known-test-dogfight-between-ai-and-human-pilot-carried-out-us-military-says-13118545>
2. Lives After Death—an example of virtual exploration for students: <https://dev.epicgames.com/documentation/en-us/uefn/lesson-plan-lives-after-death>

Part II

The Seven Visions

A Roadmap to Spatial Computing Disruption in Seven Industries

This section comprises the following chapters:

- *Chapter 4, Vision One – Transportation Automates*
- *Chapter 5, Vision Two – Virtual Worlds Appear*
- *Chapter 6, Vision Three – Augmented Manufacturing*
- *Chapter 7, Vision Four – Robot Consumers*
- *Chapter 8, Vision Five – Virtual Healthcare*
- *Chapter 9, Vision Six – Virtual Trading and Banking*
- *Chapter 10, Vision Seven – Real-Time Learning*

4

Vision One: Transportation Automates

Autonomous vehicles have become a reality, and we can now summon vehicles that will arrive without humans in them. This has deep implications for humans and the cities and infrastructure we use. They also collect data that will let us see ourselves differently and develop new services that may affect everything from how we go out to eat, to how we shop and find other services, such as nightclubs. Here, we'll dig into these deep changes and the technology that brings them to us.

A Future Vision

It is the 2030s. You are in front of Chicago's Navy Pier. It is raining, so you need a ride back to your hotel. You look into your Apple glasses and say, "Hey Siri, get me a ride."

"Okay, Mr. Smith, would you like your usual luxury ride?"

"Yes."

"A Tesla Model X will arrive within 3 minutes. Sorry, we are experiencing much heavier vehicle traffic than usual."

Within a few seconds, your glasses show you where your ride is with "A/Luxury" for "autonomous"; it is coming closer to you. You can see a map of your area and spin it with your hands in mid-air, just in case you don't like the default view. You see a 3D model of your Tesla driving toward you on a 3D map.

Why were you spinning the map? You wanted to see the city from the top down to see what the hold-up was. With this view, you can see thousands of other data sources, all showing the city in some form of 3D. There are little robots delivering things on sidewalks, showing you their surroundings. There are people walking around giving random bursts of 3D data to the map. There are tons of autonomous cars, each with a variety of sensors providing high-resolution 3D data. Finally, you can see data from cameras and 3D sensors hidden in light posts and stoplights. The city is awash with 3D bubbles of live data, which shows that there's a big wreck on the freeway. You can even see that an "M" caused the wreck—"M" for manually driven. Those cars give you data too, but it isn't 3D, just location data.

Your Spatial Computing glasses, contact lenses, or maybe a brain/computer interface someday will let you see all the data. As your car gets closer to you (this one is in the Tesla network), you can see it glowing so that you can make sure to get into the right car and not into someone else's Tesla. This visualization isn't the only thing that will change about the transportation system of the future, either. Thanks to combining databases, you will get directed from car to public transport to private drone to shared rides and back again. In the next decade, major parts of the transportation systems in many countries will be automated. So, on that ride through Chicago, you'll see trucks, trains, buses, delivery robots, and many cars roll by that have no humans driving them.

Because we have many vehicles on the road that are automated, we'll start seeing major changes to the infrastructure we drive on, too. Special lanes for automated vehicles are starting to appear. Why? Dozens of vehicles can be chained together in a convoy, which will make those lanes more efficient and reduce traffic. Already, Elon Musk is promising that Teslas will start cataloging potholes, both to report them to other drivers and to report them to crews that will fix them. His fleet added automatic lane changes so that you can see a car ahead of you finding a new pothole, and then report it to a cloud-based system that keeps track and tells all other cars, for example, to "avoid lane number one because there's a big pothole there." Even alert human drivers might catch on that all Teslas are changing lanes in an area and join them too.

Vehicles will also be controlled by the central city government in conjunction with several corporate systems that have implemented a traffic control system for watching each car get on and off freeways. It even tells those cars where to go to pick up riders, to do work such as deliver food or packages, or go for a charge or a carwash.

How do we know all this? Well, we've talked with quite a few people who are building this kind of system. But the truth is that, in 2024, the system was built and is being tested in a very small number of neighborhoods, in cities such as Pittsburgh, Phoenix, Mountain View, and San Francisco, where companies such as General Motors, Waymo, and Tesla are testing new autonomous systems.

Cars have come a long way from the slow, smelly, and loud vehicles of the early 1900s to being something that we all rely on.



Figure 4.1: Photo credit: Robert Scoble. Early image of a Google self-driving car (now branded as Waymo) on Freeway 280 near Palo Alto with an early-model Velodyne LIDAR system on top, spinning furiously away back in December of 2009.

But that reliance will soon be far deeper as transportation systems turn into utilities, where vehicles are seen less as something you own, but as something you put to work. However, all of this takes a lot of money to accomplish; several companies have run out of money in their pursuit of autonomy. Here, we cover those that are still in play, as well as some that have gone under due to lack of continued funds.

The Paths to Autonomy

If you sit in a cafe in Mountain View, or a variety of other places near San Francisco, and watch the world go by, you'll see a variety of autonomous cars roll by your table from companies such as Zoox, Waymo, and Cruise, as well as Teslas, which are more popular here than most other car brands. On the freeways nearby, you're likely to see semi-trailer trucks loaded with futuristic sensors.

It's clear from just watching the contraptions being tested on these roads (along with a few other cities, such as Phoenix, Arizona) that the future of transportation will be different than it is today, but most people are still grappling with just how different it will be. Lots of us have arguments about when autonomous vehicles will be far more ubiquitous than the Waymos that you see pass through here every few minutes.

Today, even Waymos are limited to certain neighborhoods, and there's a technical reason for that; they can only go down streets that have been imaged a few times with high-resolution cameras and lasers and have been tested several times. There is even an event called the Indy Autonomous Challenge—the only one of its kind that is a racing series for autonomous racing cars.

Autonomous cars aren't really something new. Industrialist Norman Bel Geddes created the first self-driving car in 1939. Back then, it was an electric vehicle guided by radio-controlled electromagnetic fields, which were generated with magnetized metal spikes embedded in the roadway for a "Futurama" exhibit. In 1958, General Motors took it further with a car that had sensors that could follow a wire embedded in the road. That project was aborted after the government stopped funding, even though it was shown that it would reduce accidents by 40 percent, and would also increase road capacity by 50 percent.

The idea didn't really come back into public view until Google started driving cars on Silicon Valley freeways in late 2009. Co-author Robert Scoble had the first video of that car on YouTube in December 2009 when he noticed a Prius with a weird spinning LIDAR on top (LIDARs have dozens of lasers that spin 30 times a second, making a 3D map of the road and everything around the car). He didn't realize it was a self-driving car, because there was a driver who immediately put on the brakes to get away from his camera. It wasn't until October 2010, when the New York Times wrote about the project, that people realized what those weird cars were doing on the local freeways around Mountain View.

Google's efforts actually came out of Stanford and Carnegie Mellon, but mostly Stanford, which, under Sebastian Thrun's leadership, had won the **Defense Advanced Research Projects Agency (DARPA)** Grand Challenge in 2005. That challenge was a 150-mile road course. It was a race to see if anyone could develop a car that could navigate by itself. The first time the challenge was held, no one completed the course.

What DARPA was hoping for was a new type of vehicle capable of navigating roads without needing to be manned by soldiers. They were losing tons of vehicles to crude roadside bombs. We visited Nellis Airforce Base in Las Vegas to visit soldiers who had to deal with this problem while deployed in Afghanistan, and they showed us blown-up Humvees and supply trucks and talked about fellow veterans who had been maimed or had died in these blasts. The trucks and supplies could be replaced, but the human cost was adding up.

On our own roads, the carnage was high, too. In America alone, more than 30,000 lives are lost every year to car accidents. When he was 18, Sebastian Thrun's best friend died in one of these accidents, and it drove Thrun to study how to use computers to increase safety on our nation's roads.

He told us that he didn't realize the problem would prove to be so daunting. When he started out as a Stanford student, he thought it would take a few years. Turns out it will probably take 20, and, even then, it probably won't work on every road or path in the world, nor in every condition. That could take another decade. This gets us to another path to autonomy that started around the same time: Uber. When Uber was started, it was just a year or so after the iPhone shipped in June 2007. It was partly invented in a Paris, France snowstorm. As we (co-author Robert Scoble, along with a group of other speakers at the conference) were coming out of the LeWeb conference, Travis Kalanick and Garrett Camp were talking with others about their idea. We were all stuck without cars and transportation in Paris, which had ground to a near halt due to the unusual snowstorm. Camp and Kalanick wanted to have a way to see where cars would be on their new iPhones. Within a year, the two built a company to do exactly that and Uber was born. Today, Uber is valued more than most car companies, which shows that investors are behind the idea, and it has major competitors in Didi in China and Lyft in the US, both of which also are valued at tens of billions of US dollars.

One reason that so many billions of dollars are being invested in autonomous cars is that the economics will dramatically change Uber and Lyft, because as you remove humans from driving, you radically change the potential economics and you get a far more consistent user experience, which is key to building a great brand. Riccardo Giraldi, former Senior Director of Product Experience for Zoox (a company that is building its own taxi vehicle from scratch based on autonomous driving), asked us to think about Starbucks. Starbucks is a great brand, not because of its amazing coffee, but because it's extremely consistent, ubiquitous, affordable, and personalized. All of those qualities, Giraldi told us, will be possible once autonomous cars are rolling in enough cities.

Since Waymo started in its first one or two cities in 2019, we expect that it will take just a few years for automatic cars to be operating in 20 cities (currently there are 10 cities where Waymo can operate), and in 100 cities by the end of the 2020s.



Figure 4.2: Photo credit: Robert Scoble. A self-driving car in Velodyne's parking lot in Morgan Hill, California. Velodyne makes the LIDARs (the spinning set of lasers and computers on top of the car) for many of the companies doing self-driving research.

This race has already created a new ecosystem of sensor development and companies that make machine learning systems and test suites. One of the more interesting people to meet in this area is David Hall. His company, Velodyne, sold the best subwoofers to audio enthusiasts back in the 1980s. As he moved his production to China in the 1990s, he found he still had a group of smart electro-mechanical engineers and a huge warehouse in Morgan Hill, California, which is in the southern part of Silicon Valley. He heard about the DARPA challenges and tried putting together a car to compete in it, but he quickly learned that he would be outspent by universities such as Stanford and Carnegie Mellon, and so his entrepreneurial skills went into action.

Instead of building his own car, he decided to build the sensors that the other teams would need. His Velodyne LIDAR was the spinning thing on top of that first Prius that Google was driving around Silicon Valley, and today, Velodyne is moving to new semiconductor LIDARs that no longer spin and can be hidden inside car bumpers.

Falling Car Transportation Costs

Transportation is about to get dramatically less expensive. Insiders from Sebastian Thrun, who ran the first Google autonomous vehicle team before he went on to run Udacity and an electronic flying vehicle start-up called Kitty Hawk, to Elon Musk, CEO of Tesla, give reasons why transportation will get cheaper. The biggest among them is that cars will no longer sit in garages most of the day but will go out and “work” on your behalf. There are other economic advantages, too. The original Stanford team told us that autonomous vehicles could save pennies while they drive because they would be more efficient than human-driven ones, even going up and down hills more efficiently, or joining caravans on freeways.

If you buy a car today, you will probably buy a car that costs thousands to tens of thousands of dollars. You’ll also be paying to maintain it and to keep it insured. And with all that expense, you’ll only use it an hour or two a day; the rest of the time, it will sit in your garage. That’s hardly a great economic decision. Imagine spending capital on a factory and only running it for a couple of hours a day. That doesn’t make sense, right?

Car-sharing networks like Uber and Lyft have shown that a car can be used much more frequently, which shares the cost among many riders, but there are much deeper cost advantages coming due to electrification and the development of autonomous vehicles.

If you take the cost of buying a \$70,000 car, and all the other costs, like insurance, fuel, maintenance costs of tires, filter changes, windshield wipers, oil changes, and brakes, it comes out to around \$4,000 a month if you drive it around the clock. It’s cheaper if you buy a \$35,000 Toyota.

Now, divide that by the 733 hours in a 31-day month. That comes out to a little less than \$5.50 an hour. Add \$15 minimum wage and your wholesale costs come to just a little over \$20 an hour. Uber or Lyft turn that into \$60 retail. (The standard is to always triple wholesale prices to arrive at retail. You could easily check this if you so desired: order an Uber and ask your driver to drive you around for an hour. The bill will be more than \$60).

If you could rent the same car out without a driver, your costs would only be \$5.50 an hour wholesale, or about \$17 retail. \$17 versus more than \$60. That is the economic power of driverless cars. It explains why so many billions of dollars are being invested into companies like Cruise (General Motors), Waymo (Alphabet (Google)), Zoox, Aurora, Tesla, and Daimler (Mercedes-Benz), not to mention many others, including Lyft and Uber themselves.

So, when self-driving hits, really hits, the economics will totally change. This change will be as big of a deal as when Henry Ford turned on the first assembly line and made cars available to a much larger percentage of Americans due to dramatically lowering the costs of buying a car.

When the economics change, so does the utility. As transportation gets cheaper, we will use it for more. If we could sleep in our cars as they drive, we could live further from our jobs. Our cars could do work while we sleep, or work. “Hey Tesla, pick up my laundry,” we might soon say.

What does this have to do with Spatial Computing? Everything. First, the technology underneath a self-driving vehicle is very similar to the technology that makes **Augmented Reality (AR)** glasses possible. Eight cameras, sensors for gravity and movement, and 3D sensors—all joined together, and AI does the work. These glasses will also be used to both control transportation and to enjoy it more due to virtualized screens and 3D social networks, games, entertainment, and education, which we'll go into elsewhere in this book.

There are already companies doing that, like Holoride, which spun out of Audi. Holoride won best of the Consumer Electronics Show in Las Vegas in January 2019 and gives passengers a way to play VR games. As cars convert to electric, more minutes will be spent charging, and many more families will take VR headsets on long drives to make those charging stops more fun. As companies like Apple and Facebook introduce AR glasses, there will be much more demand for the kinds of games that Holoride had planned. In 2024, Holoride filed for bankruptcy due to not being able to raise more funds.

The co-founder of Holoride, Marcus Kuehne, had told us they had lots of plans to expand as more Spatial Computing devices like AR glasses get introduced and as cars start to become more automated. Already, if you have a Tesla, you don't have to pay much attention as your car drives on freeways. Soon, people who have these bleeding-edge cars will want to do a variety of things while on the freeway, like watching movies or TV, answering emails, looking at social media or news, or doing deeper work like catching up on a PowerPoint presentation or a spreadsheet before getting into work.

This demand goes way beyond games. We expect Spatial Computing glasses to bring all sorts of utility to driving, being a passenger, getting deliveries, and being a passenger in a self-driving vehicle. Not to mention that lots of insiders say that the technology that will run the Apple glasses is very similar to the technology that lets a computer system drive a car (all current XR headsets already have these Spatial Computing abilities). In many ways, they will overlap in the next decade, one adding to the other.

These same glasses will be used to fill in gaps in autonomy, too. Sometimes, autonomous vehicles may get stuck and will have to be handed over to remote drivers. Those remote drivers will be wearing Spatial Computing glasses to see what the car is seeing.

Even if you “own” your own vehicle, which many will still need—for example, if you are a mechanic who needs your tools to always be in the back of your pickup truck—the costs will come way down and you’ll pay a per-hour and per-mile fee that will be lower than what you are paying today. Why? Electric vehicles are cheaper. Autonomous vehicles are too, and you’ll be able to do more as you use them, since you’ll be able to use your Spatial Computing glasses to watch movies, play games with your friends, or simply catch up on some sleep while your truck drives you to your next job. All of this dramatically lowers both the real costs of using transportation and will bring deep changes to every part of life.

Autonomous Vehicle Round-up

Each **autonomous vehicle (AV)** has a slightly different configuration of equipment and systems needed to make things autonomous, so we will review some of the main setups to give you an understanding of the technology. Public information regarding setups is extremely variable due to competitive concerns.

In addition, we will touch upon autonomous trucks and autonomous flying cars—officially called autonomous eVTOLs, which stands for “electric vertical takeoff and landing” vehicles.

Before going into companies’ autonomous configurations and other information, we’ll pause to explain exactly what we mean by “autonomous.” The **Society of Automotive Engineers (SAE)** International is a global association of more than 128,000 engineers and related technical experts in the aerospace, automotive, and commercial vehicle industries, and it sets the standards for the five distinct levels of autonomous technology in vehicles:

- **Level 5** is full driverless automation, where human control will never be required and, in fact, may be outlawed on public roads.

- **Level 4** allows full automation in most situations, within a specific geographic area and during appropriate weather conditions, needing driver intervention infrequently.
- **Level 3** allows full automation in certain situations, notably freeway driving, with driver intervention needed frequently.
- **Level 2** offers a combination of autonomous functions, such as basic lane-keeping autonomous steering, in combination with active cruise control; passengers actively driving is the norm.
- **Level 1** includes some very basic autonomous functionality, including electronic stability control and active cruise control technologies, with drivers expected to maintain complete control.

Most autonomous vehicle companies are currently operating on Level 4 autonomy, most notably on known scoped and mapped-out roads and testing tracks. Tesla is currently operating at Level 2. Level 4 and above is typically marked as autonomous.

Now that we've covered the generally accepted meaning of autonomous, let's discuss some of the companies that are working on these autonomous vehicles.

Zoox

The main reasons for the existence of autonomous vehicles include safety, mobility, and sustainability. Looking at the US-based infographic from Zoox below, you can see that 94 percent of car crashes have been attributed to human error, many millions of adults have disabilities that preclude them from driving freely, and massive amounts of greenhouse gasses are emitted by passenger vehicles, with much of that occurring during rush hour traffic.



Figure 4.3: Credit: Zoox. Statistics that can be related to non-autonomous passenger vehicles.

Zoox is one of the companies that is pioneering the autonomous car revolution. It is one of just a few that is actually building the car from the ground up, rather than taking an existing car model and adapting it to an autonomous system. We talked to Mark Rosekind, former Chief Safety Innovation Officer of Zoox, and he told us, as an example of improving safety, Zoox builds the back seats of their vehicles to have the same form features as the front seats, which are safer by design. Additionally, Zoox cars are not built with a steering wheel. Steering wheels cause many injuries during car accidents.

After it was founded in 2014, Zoox was acquired by Amazon in 2020. In addition to building its own cars, Zoox has developed its own autonomous-driving software, electric vehicle program, and urban ride-hailing system. In March 2024, Zoox expanded its autonomous driverless testing near its headquarters in Foster City, California, and in Las Vegas, operating along five miles of road from the south end of the Las Vegas strip. As of June 2024, Zoox will be opening Las Vegas to the public. Zoox will become the first company in the US to operate a purpose-built robo-taxi on public roads at relatively high speeds, thanks to its deployment in Las Vegas. Zoox recently announced plans to start testing its driverless robo-taxi rides in Miami, Florida and Austin, Texas.



Figure 4.4 Credit: Zoox. The latest Zoox autonomous driverless car.

Zoox autonomous cars come equipped with multiple cameras, four LIDAR systems, radar (radio detection and ranging) and proprietary sensors, along with GPS and proprietary mapping data systems. LIDAR uses light in the form of a pulsated laser to measure distances and is the most expensive sensor system of the sensors suite used, typically upward of 10 times more expensive than radar.

The LIDAR units can be seen around the top of the car, along with radar and proprietary sensors and cameras. Radar units can be found near the front and back bumpers.

For all current autonomous vehicle companies, the cameras and sensor suite work together with computer vision software using machine learning in order to perceive surroundings and environments and predict future actions of road objects, in order to plan the car's movement accordingly.

All current autonomous vehicle companies navigate roads by way of a combination of virtual computer simulation testing and modeling, which uses machine learning and prior meticulous mapping of the roads to be traveled on. All autonomous vehicles currently use LIDAR. Tesla, whose cars are not yet autonomous, does not utilize LIDAR as a part of its navigation systems (more on this later).

Ford

Ford was on the fourth generation of its self-driving test vehicle, using Fusion Hybrid and Escape Hybrid sedans for their self-driving test vehicles that were on public roads in the cities of Miami, Pittsburgh, Dearborn, and Palo Alto, and had plans to branch out to Austin, Detroit, and Washington D.C., when it decided in 2023 to exit its Level 4 strategy.



Figure 4.5: Credit: Ford. Austin was chosen as the third city, behind Miami and Washington, DC, where Ford had announced that they would launch autonomous vehicles by 2021.

Ford had partnered with Pittsburgh-based Argo AI, an Artificial Intelligence and robotics company, with a \$1 billion investment to develop what Ford calls the “Virtual Driver System” for its self-driving vehicles. In 2019, Ford announced that it was teaming up with Volkswagen to build a system for a Level 4 vehicle, with Volkswagen putting \$2.6 billion of investment into Argo AI. Argo AI shuttered in 2022 after failing to secure further investment.

General Motors' Cruise

With regard to General Motors, their autonomous vehicle company is called Cruise. Each Cruise vehicle is equipped with five LIDAR units, 16 cameras, and 21 radar units and calculates its path on the road 10 times per second.

General Motors acquired Cruise in 2016 for reportedly over \$500 million, and it has grown from around 40 employees to over 200 employees.



Figure 4.6: Photo credit: Cruise. A General Motors' Cruise self-driving car navigates the streets of San Francisco with a safety driver behind the wheel.

Cruise had suspended operations in the fall of 2023 after a San Francisco accident, but it announced in April 2024 that it would gradually return to city streets, starting with Phoenix.

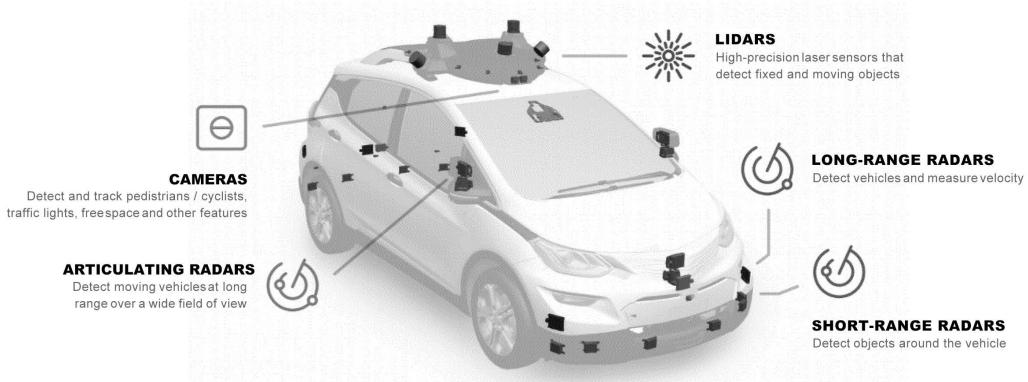


Figure 4.7: Credit: Cruise. Cruise's autonomous car system has more radar units than other companies systems.

Waymo

Waymo, the autonomous vehicle business under Alphabet (Google's parent company), has been saying for a while that its self-driving system is designed to operate without a human driver. In October 2019, Waymo sent an email to customers of its robo-taxi service in the Phoenix area that indicated, "Completely driverless Waymo cars are on the way." In March 2024, it was announced that Waymo would be bringing their robo-taxi service to Los Angeles in that month and also to Austin by the end of the year. It also now operates in parts of San Francisco and Phoenix. It is testing its technology in Atlanta. Uber uses Waymo cars as part of its service, employing it for Uber Eats as well, in the Phoenix metropolitan area.

Waymo's systems include three types of LIDAR that were developed in-house—a short-range LIDAR that provides an uninterrupted view directly around it, a high-resolution mid-range LIDAR, and a powerful new-generation long-range LIDAR. Other parts include a camera and Computer Vision system, radar, and self-driving software that incorporates machine learning trained on more than five million miles of on-road experience in six US states, more than 25 cities, and more than five billion miles of simulated driving. Waymo first began testing its technology in 2009 in and around its Mountain View headquarters.

Waymo has agreements to develop self-driving vehicles and services with Renault and its Japanese partner, Nissan Motor, and has deals to use minivans made by Fiat Chrysler Automobiles NV and Jaguar iPace sport-utility vehicles in its robo-taxi fleets.



Figure 4.8: Photo credit: Waymo. A portrait of a Waymo autonomous car.

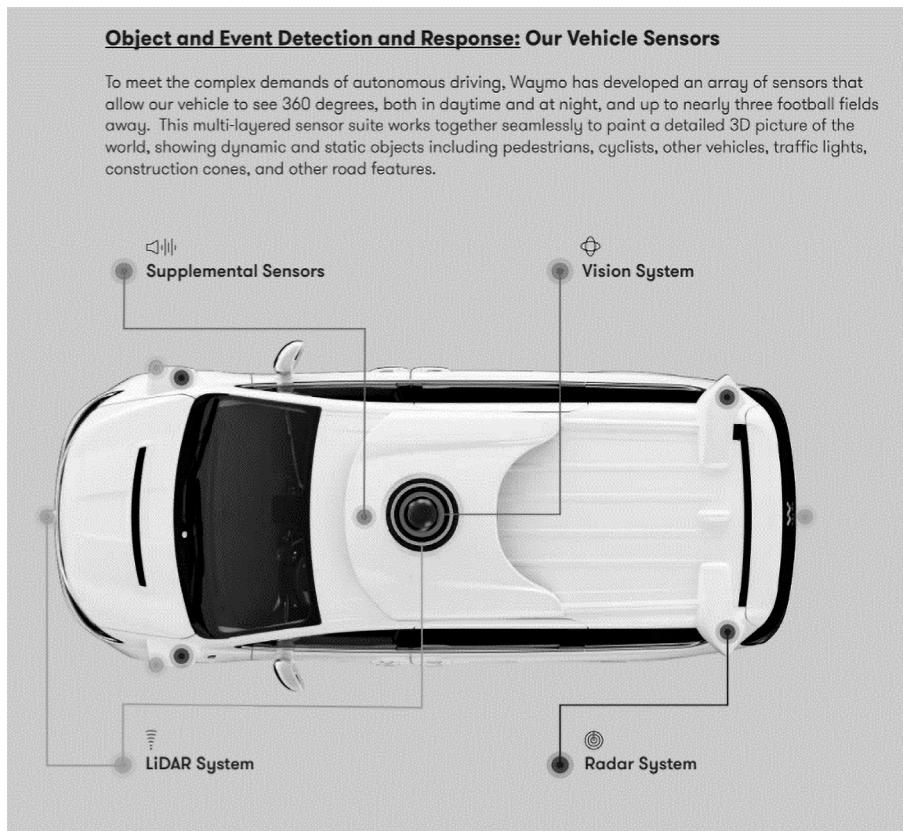


Figure 4.9: Credit: Waymo. Details of Waymo's vehicle sensors.

Tesla

We are including Tesla in this section, even though it is currently only at Level 2 autonomy. The reason we are doing so is to make it more clear what Tesla's current status is and what the possibilities are with regard to fully autonomous driving.

In comparison to others who use LIDAR systems, which can be 10 times more expensive than using radar, Tesla does not use LIDAR, but rather uses just radar, cameras, machine learning, ultrasonic sensors, and a custom computer system with a custom chip designed to enable full self-driving capabilities in the future. The next-generation chip is said by Elon Musk to be three times better than the current chip and about two years away.

The current system has a forward-looking radar, eight cameras, 360-degree ultrasonic radar, and GPS.

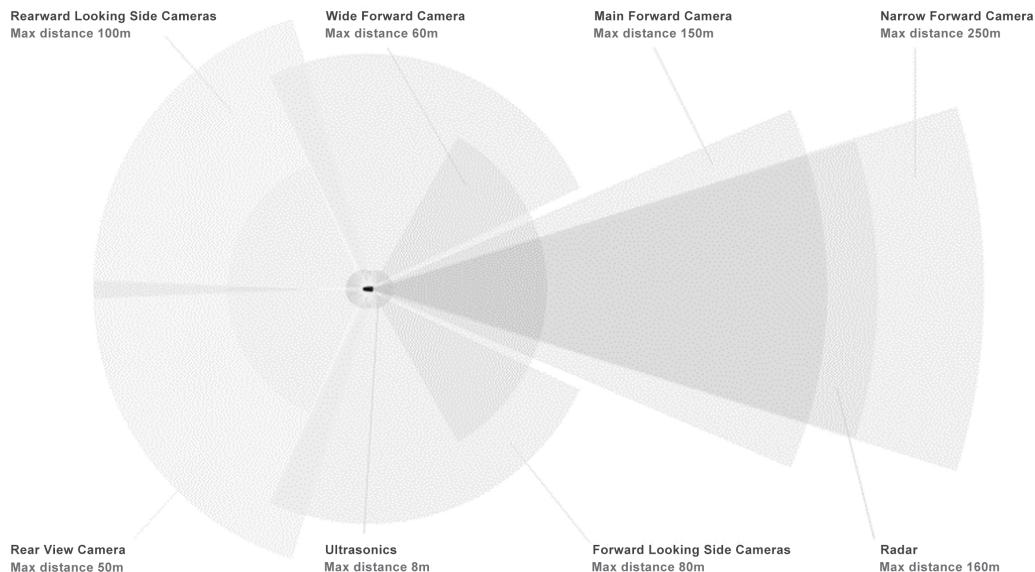


Figure 4.10: Credit: Tesla. Tesla car camera and sensor system distances.

Musk believes that it is possible for autonomous vehicles to rely solely on radar, cameras, ultrasonic sensors, and GPS, rather than having to include LIDAR, which is expensive. And there are some key people that we have interviewed that also believe this to be the case, but they have asked us not to reference them or quote them.

Right now, there are only a few features that are part of Tesla's Level 2 automation: traffic-aware cruise control, autosteer, auto lane change, navigate on autopilot, autopark, summon and smart summon, and traffic light and stop sign control. Two terms, summon and smart summon, need more explanation. With summon, the car can move in and out of a tight space using the mobile app or key, and with smart summon, the car can navigate more complex environments and parking spaces—maneuvering around objects as necessary to come and find you in a parking lot.

Apple

In February 2024, Apple closed down its formerly secret endeavor, "Project Titan," whose product was an autonomous car. This autonomous car was also to be an **electric vehicle (EV)**. Apple had worked on this project for nearly a decade with billions spent. The reason for the shutdown is thought to be due to too many obstacles, objectives (okay, it's an autonomous car and it's also an EV), and a lack of unified direction from within the company.



Figure 4.11: Photo credit: Andrej Sokolow. One of Apple's "Project Titan" cars.

Hyundai, Pony.ai, Aptiv, and Aurora

In 2020, Hyundai launched a pilot ridesharing service using prototype autonomous cars. A fleet of autonomous Hyundai Kona Electric prototypes started to provide free rides around Irvine, California.

Hyundai had partnered up with Pony.ai, a Chinese start-up with Toyota ties, on the autonomous-driving system utilized in the rideshare cars. Pony.ai provided Hyundai with the sensors and proprietary software that enables cars to analyze data and make driving decisions.

BotRide, the app that will be used to hail these cars, was co-developed with Via, a ridesharing company.



Figure 4.12: Photo Credit: Hyundai. Hyundai had teamed up with Pony.ai, Aptiv, and Aurora.

In 2020, Hyundai announced a \$4 billion joint venture entitled Motional with Aptiv to develop Level 4 and Level 5 technology. Aptiv exited from the joint venture in February 2024. Motional still operates a robo-taxi service in Las Vegas on the Uber, Lyft, and Via platforms, with human safety operators still behind the wheel. Additionally, Hyundai had taken a minority stake (undisclosed amount) in the start-up Aurora, but in 2022 Aurora decided to focus on autonomous trucks.

Autonomous Trucks

We've been told that **autonomous trucks** (ATs) could be 1.5 times more expensive than conventional trucks due to the higher initial investment, and there has been a lot of talk about truckers losing their jobs due to ATs.

However, there is currently a massive driver shortage. As truckers retire, they are not being replaced. Younger generations are going for less demanding careers in other industries.

Companies that operate ATs, and those that provide software for ATs, will be able to recoup the relatively high initial investment needed to enter their markets in a number of ways. ATs are not subject to hours-of-service regulations, which limit the amount of time a driver spends behind the wheel. With ATs, driving hours could be increased from 11 hours per day to 20. Trips will be faster due to more proficient driving due to optimized routing software and the elimination of the need for rest periods.

In terms of the kinds of change the trucking industry will undergo due to ATs, there will be a number of stages. The first two stages will feature “platooning,” a technique used to wirelessly connect a convoy of trucks to a lead truck, which allows them to operate much more closely together safely and increase fuel efficiency. The first stage will be platooning with drivers at Level 3 autonomy. It will take three to five years to develop the networks of connected convoys.



Figure 4.13: Photo credit: Starsky Robotics. An example of a Starsky Robotics autonomous truck.

The second stage, coming in about five to seven years, will feature driverless platooning. While on an interstate, a driver will be in a lead truck and unmanned trucks will follow close behind. Drivers will take control of each vehicle once it leaves an interstate.

In about seven to ten years for the third stage, we predict that ATs will operate under Level 4 autonomy. Trucks will be unmanned on interstates and other “geofenced” areas without a platoon. Once off the interstate, drivers will take the ATs to their ultimate destination.

The first fully autonomous trucks, operating at Level 5, should appear in more than 10 years. This is the fourth stage, and over a long period of time, there will be a general move for a preference toward Level 5 ATs.

We are predicting the stages and timing for ATs, but not for AVs due to the complexities of city streets and political/legal concerns.

There are a number of companies that are working in the AT space, including Ike, Embark, TuSimple, Kodiak Robotics, and Aurora. Waymo pulled out of the trucking business in 2023. Starsky Robotics shut down in March 2020 due to operational issues and a failure to get more investment.

Flying Cars (eVTOLs)

In 2020, we talked with Sebastian Thrun, then CEO of Kitty Hawk Corporation and chairman and co-founder of Udacity, a for-profit online educational organization. Prior to this, at Google, Thrun founded Google X, home to projects like the Google self-driving car (now Waymo) and Google Glass, and at Stanford University, he was a research professor in the Computer Science department and director of the Stanford Artificial Intelligence Lab.



Figure 4.14: Photo credit: Kitty Hawk. One of Kitty Hawk's eVTOL vehicle types, Heaviside, navigating the skies.

Thrun was very passionate about eVTOLs, and he believed that there was a great chance that operational autonomous eVTOLs would come before autonomous vehicles. The main reasons for this included the fact that airspace is much less cluttered than the roads on the ground, and that the technology for monitoring airspace had already been developed. Since eVTOLs travel in a straight line in the air rather than having to follow non-direct and non-linear roads, they would be more energy-efficient than AVs.

Additionally, Thrun believed that it was possible that eVTOLs could navigate quite well without needing LIDAR, with a nod to the possibility of AVs not needing LIDAR due to the great strides the combination of machine learning, Computer Vision, and camera technologies have made.

Two of Kitty Hawk's eVTOL vehicle types, Heaviside and Cora, had been primed to become autonomous, with Boeing partnering up on Cora. A feature of Heaviside was technology that made it roughly 100 times quieter than a regular helicopter. Being quiet is very important when it comes to taking off, navigating, and landing in urban areas, where vertical take-off and landing helipads, which are being called "Skyports," will be necessary. Both Heaviside and Cora were built to go no more than a little over 100 miles on a charge, with 70 miles being an average, and travel at 180 miles per hour, being about 10 times faster than travel on the ground by road. Charging was a breeze, taking no more than 15 minutes for a full charge, with the batteries being light and unremovable. Unfortunately, Kitty Hawk shut down operations in 2022 due to money concerns.

Some other notable companies that are developing eVTOLs, and then autonomous eVTOLs, are Joby, which bought Uber Air (which would enable four-person ridesharing flights in densely populated areas), Boeing's Aurora Flight Sciences (which Uber Air had partnered up with), Archer, Wisk, Pivotal, Alef Aeronautics, Chinese company EHang, German company Lilium, Singapore company Eve Air Mobility, and German start-up Volocopter GmbH, backed by Daimler AG and the Mercedes-Benz maker's biggest shareholder, the Zhejiang Geely Holding Group. Boeing CEO Dennis Muilenburg has stated that he believes that commercial flying taxis could be a reality by 2025. In fact, Volocopter announced that their VoloCity air taxis will be available in Paris in 2024, being the first to offer electric air taxis in Europe.



Figure 4.15: Photo credit: Boeing. Boeing's first test of its Uber Air air taxi in January 2019.

Alef Aeronautics is developing a car that you can drive and fly and plans to start production in 2025. The car known as the Model A will cost \$300,000. Lilium, which is backed by the Chinese company Tencent, is planning to offer its jet in 2026 at a cost of \$9 million and is doing test flights in Southern Spain. The cost is high because it can hold several seats and can go on longer regional flights. EHANG has a production certificate to begin producing eVTOL aircraft, at a price of around \$400,000.



Figure 4.16: Photo credit: Alef Aeronautics. A mockup of its flying car.

Before any cultural or technological shift, particularly one as big as Spatial Computing looks to be, there's lots of arguing and resistance. Next, we will consider the extent of this "culture shock," what some of the arguments presented by skeptics might be, and how we foresee these arguments being overcome.

The Autonomous and Electric Culture Shock

When it comes to transportation shifts, the arguments are varied and deep. Bring up electronic cars at literally any cocktail party and you'll hear a spectrum of reactions from "what a stupid idea—you have to charge for 40 minutes" to "people who stick with gas are destroying the earth."



Figure 4.17: Photo credit: Robert Scoble. New meets old: this book's co-author's family drove one of the first 10,000 Tesla Model 3s built to Detroit, where they visited Ford's first plant. Henry Ford, in this building, changed cars by building the first assembly line, which dramatically reduced the cost of buying one.

Add talk about autonomous vehicles and you might experience an all-out brawl: “I love to drive.”

The thing is, our culture is already changing due to early experiences, and while the resistance is deep, we look back to prior changes and, every time there is a big shift, we hear the same kind of resistance. We look back at when the automobile first appeared in society, and there were many who didn’t like the contraptions, either. They scared the horses, for one, so much so that early owners had to have someone walk in front of the car to warn people one was about to come down the street.

How many people told us “I’ll never carry a phone everywhere?” or “I’ll never use social media”? As we have written previous books about technology shifts and were among the first to try many a new technology, we can answer this for you: many.

That isn’t to say that the resistance can be written off wholly. This is not true, particularly when it comes to transportation. Their views, as expressed by especially conservative lawmakers, could bring a variety of legislative approaches. Already, we are seeing a vast difference in the technologies used on the streets. Some cities—London, for example—have blocked ride sharing, like Uber or Lyft, and still do. Others already have self-driving vehicles being tested out. In Phoenix, that brought tragedy, as an Uber car hit and killed a pedestrian in the early hours of one morning.

Teslas have been blamed for a variety of deaths, including an Apple engineer who was tragically autopiloted into a barrier on a freeway offramp in Silicon Valley.

You might think that the federal government will decide, at least in the United States. That's not true, however, if you look at the vehicle code: each state decides on laws. This is the case even for interstate highways; states determine the laws upon any stretch of highway passing through said state. It's unclear how government administrations will work together to regulate autonomous vehicles, and there are many unanswered questions, from responsibility for deaths and wrecks to whether or not you will be allowed to consume alcoholic drinks in an autonomous vehicle (today, that's not legal, at least in California, or if you are caught even in the back seat of a vehicle that's moving and you are intoxicated), you could be arrested for driving while intoxicated. Protestations of, "but hey, my Tesla is driving," have not yet been accepted by courts. We'll see many of these legal questions start to get settled in the 2020s.

You see this culture shift under way as people drive huge pickup trucks with horrible fuel mileage driving next to electronic cars from Nissan, Tesla, Chevy, and others, and you see the shifts on social media as people argue about which automaker produces the best product. While over the last few decades those arguments might have been between Chevy and Ford or American versus Japanese or versus European, the arguments in the 2020s will be about technology ecosystems, Spatial Computing apps, and who has the best AI and on which roads.

Already, owners of Teslas have been captured on video sleeping on the freeways while their cars drive. The fact that these videos have gone viral and are discussed on TV news shows that a new usage pattern has started to arrive, and it's proving to be unnerving for many. There's a feeling that it's unsafe for everyone. Even owners of Teslas derided the first sleepers, but we posit that sight will be extremely common by the end of the 2020s, and it won't be done in just Teslas that only fairly well-off people can afford.

The next weird thing to get used to will be cars driving around without any humans in them at all. Tesla turned on the ability for cars to move around shopping mall parking lots this way, which generated another round of video derision. However, in 2024, an owner was always nearby, just in case something went wrong. Sometime soon, the owner will no longer be there, and our culture will continue to shift.

For businesspeople, these shifts bring new fears. Soft drinks and snacks are often consumed after a quick visit to a gas station, particularly on long trips. Consumer goods companies are already planning to see sales radically change as those distribution channels disappear and new ones appear.

On our trip around the United States, we found that many Tesla supercharging stations were near restaurants, or sometimes even inside existing gas stations, so new patterns of traffic are springing up. When freeways were first brought through communities, similar changes also appeared. If there was an off-ramp, real estate values frequently went up and businesses appeared. The new off-ramp might include many charging stations or services for drivers on longer trips. Tesla already has huge stations with coffee, restrooms, and other services available between cities like San Francisco and Los Angeles, which didn't exist before people started using its electric cars to travel those multi-hundred-mile distances.

Electric vehicles also can also be used as micro-offices, or even temporary homes. Lots of owners talk about using their Teslas as camping vehicles, asking how to keep the AC system running properly all night long, which is something that would seem dangerous in a gas-powered vehicle, or at least destructive to engines and the environment. Now that electric vehicles run more like office buildings, you can keep them temperature-controlled and can work in a hot or cold parking lot while waiting. The silence from electric engines lets you do a variety of things, from videoconferencing to playing games on your phone or laptop. Already, a sizeable industry has sprung up with accessories for that kind of usage, and we know several businesses that have waiting lists for upgrades to cars such as tinting windows, which makes them nicer to use as portable offices.

Quite a few automakers, including new ones like Zoox, are seeing that trend and have built new vehicles that will not have a steering wheel, but rather chairs and tables for having a business meeting or maybe a family meal inside as you are being driven around.

Only a few, like Brad Templeton—who was an engineer on Google's early autonomous efforts, and today teaches executives about technology shifts that are coming at Singularity University—have an idea of just how deeply American culture and the landscape itself will change due to the three forces of autonomous vehicles, electrification, and Spatial Computing glasses and devices. Since you will be able to sleep in these cars by the end of the 2020s, you'll drive overnight.

Imagine leaving San Francisco at 9 p.m., sleeping in your car, and arriving in the morning in Las Vegas or at midday in Salt Lake City. Those trips are uncomfortable to do by yourself, and quite dangerous due to the real possibility you'll fall asleep while on the boring straight highways that can numb even the most active of brains.

Trucks will also go through this transition starting in the 2020s. Yeah, it might take another decade after that to see autonomy hit literally every vehicle, but we can already predict the decimation of many a roadside diner or hotel focused on the more than a million truckers who drive in the United States.

On the other hand, those diners and hotels have seen a steady stream of post-COVID-19 restrictions of new tourists who now see driving through the United States as a sort of amusement park (the same for Europe, and many other places where the drives are long and there's lots of things to see along the way).

As these drivers approach cities, they will probably need to use a service to request permission to come into the city. Brad Templeton says that more and more roads will be metered. Why? Because as we automate transportation, more cars will be on the roads, particularly near city centers, because they will be asked to do more work than before, with people possibly asking things like, "Hey Siri, can you have my car pick up my laundry?"

That means more traffic and cities will have an impulse to both regulate that traffic and attempt to profit from it. In New York City, there was a 2024 proposed regulation that employed a cordon pricing system to charge drivers to access Manhattan south of 60th Street. We predict that soon, cities will force drivers to get apps on their phones that will give them a pass to get on freeways into cities. Such a thing wasn't possible until literally every driver got a mobile phone, which brought us not only Uber, but new mapping services like Waze. Uber already works sort of like that, by the way. It tells its drivers where to go, all with a centralized control system. That system knows where every driver and passenger is, and charges accordingly. We won't be shocked to hear that Uber is called upon to help build such a system for cities. Your glasses by 2030 might have you waiting outside of a city for a "slot" for you to drive in, just like airplanes are often held on the ground if there's a problem at a distant airport. This is yet another reason to have charging and other services available in stations near freeways.

You'll be seeing more acceptance of ride-sharing, particularly between cities. In 2018, we hitched a ride between Munich and Freiberg, Germany, in the back of an Audi thanks to BlaBlaCar, a service that sells seats in everyday people's cars as they commute between cities. Such a thing was impossible to implement before ubiquitous mobile phones arrived, and as Spatial Computing glasses gain acceptance in the 2020s, these kinds of things will become much easier to use.

We envision walking down a street and seeing tons of transportation options, when you want them, all displayed by your glasses—from scooters to flying vehicles, all overlaid in 3D around you so you can see what your options are. In many ways, this represents the purification of the vision that Garrett Camp and Travis Kalanick told us about as they were thinking about creating Uber in a Paris snowstorm. We were stuck due to that snowstorm and no cars were coming to pick us up because traffic was at a standstill. Camp asked us, "Why can't we just see where the cars are?"

In the future, you might use your AR glasses while sitting in traffic to get out of the cold and have a place for you to do a video call, which of course will be radically different in the future due to Spatial Computing. No longer will we be stuck to 2D screens. Your business meeting could have virtual whiteboards, desks with virtual items on them, plus lots of 3D designs. Think about being an architect—you'll be able to work on your building design while sitting in a car that's driving you across Paris. Even if you are stuck in traffic, you won't care nearly as much as you do today because the vehicle you are in will be much more comfortable and quieter, and you'll have hyper-fast bandwidth available to your AR glasses.



Figure 4.18: Photo credit: Robert Scoble. Google Maps shows you an augmented map that helps you navigate cities better. This image was made in late 2019 in Rome, Italy. In the future, this augmentation, which uses Spatial Computing technology, will be expanded to be far more helpful and detailed.

What other things will shift about our culture in the 2020s? Where you buy your car, if you buy one at all. Let's say you need a new pickup truck for your plumbing business. Today, you need to go into a Ford dealership to get your car. Well, what if your truck could drive to you? Doesn't that get rid of the need for dealing with car salespeople? Yes. If we owned car dealerships, we would spend the next decade planning for this new way of buying a car or truck to hit our business hard. But it won't happen in one day—it'll be something you'll see coming for years. Just remember, these technologies are exponential. If you wait too long to get a different business model, it will be a lot tougher, particularly if you are the last one to make the move. The deeper issue for car salespeople will be the psychological move from needing to own a car or truck at all.

A system like Uber (or Tesla's planned RoboTaxi network) could dramatically change vehicle buying patterns and, in fact, due to the huge gap in economics, we foresee a huge problem coming for car dealerships and salespeople by the end of the 2020s.

It is not only car dealerships that will be undergoing change, but also major parts of the cityscape.

The Future City

To get an idea of how deeply cities can change in a decade, we walked around old Barcelona, Spain in November 2019. Every store accepted Apple Pay. Taxis worked like Uber. We were guided by AR on Google Maps down streets that were hundreds of years old. Apple has a store here. Every restaurant had tons of reviews and you could get reservations at many online. None of that existed in 2010.

Some cities, like Rome or Milan, will be harder to automate because of their lack of painted lines on streets, their much more chaotic pedestrian traffic, and unusual parking and traffic arrangements. Expect even these cities to see big improvements due to AR glasses and the data collection from those and autonomous vehicles, however.

We asked many people who are building these technologies how cities would change in the future due to Spatial Computing technologies like robots, autonomous vehicles, and AR glasses.

For one, as vehicles start to automate, the need for them to park and sit around unused for hours will go away. Templeton noted that about 20 percent of real estate in cities is for parking. We can see a world in, say, 40 years, where all that parking has been turned into either parks or housing to make living in cities more affordable.

The biggest change coming will be the 3D-ization of cities. Similar to what is being done inside factories, cities soon will have a "digital twin" or a digital copy. These digital twins will provide the virtual infrastructure for a whole range of new utilities and services that will be exposed in Spatial Computing/AR glasses.

We are already seeing early evidence of this as Google Maps is adding AR navigation and street signs onto city streets in many cities around the world (we tried it in Barcelona and Rome, for instance, and it worked well).

Today, maps still remain fairly static, though, and don't represent crowds or changing street conditions well. If a construction crew blocks off a street, people at home have no idea until they try to go down that street and find it blocked off. Soon, maps will become real-time and show things like that in far greater detail and with far less latency, or delay.

These maps will be greatly helped by self-driving vehicles, which continuously map out streets in 3D as they drive down them or, at a minimum, notice differences between a pre-scanned map and what actually is on the street.

The people building self-driving cars even could imagine that as cars drive down a street, they would refresh their data to include parking spaces available, what potholes are forming, construction underway, and details on accidents or crowds, among other things that would be pushed immediately to maps, not to mention the various agencies that could keep track of how to make roads better for citizens or, even, take note of new graffiti appearing or cars being left in one spot too long.

Particularly as 5G arrives in cities, both autonomous vehicles and this kind of “live mapping” will be seen as much more commonplace. Marketers might even develop a new kind of SEO. This extra data will also be used by both public authorities and transportation services like Uber, Didi, or Lyft, to move resources to places where there will be more demand. For instance, a nightclub with a line outside at 10 p.m. will probably need more vehicles nearby at 2 a.m. to give people rides home.

This new digital twin will be joined by other deep changes to cities brought by the combination of Spatial Computing glasses that users talk to, use their eyes or hands to control, and that present new visual information about the city around them.

Let’s pop around the city to see some of the places where we predict we’ll see the biggest changes. For the remainder of this chapter, we will list some of the major changes that we foresee occurring within and around our cities over the next decade.

Auto Dealers

Let’s start on Auto Dealer Row. You know, the place you go to buy cars. By 2030, we expect dealerships to start to shrink as more and more new cars come online that can be bought either by phone or by Spatial Computing glasses. Why will this happen? Cars by 2030 will be able to deliver themselves directly to users. Keep in mind that not all car sales will go away. Families who have car seats will want a “long-term car” for what the car services are calling them—the same for mechanics or plumbers who want to keep their tools in the back of a pickup truck and paint logos onto them. Either way, though, the car salesperson sure will have their jobs threatened.

Smart Roads

That new car will communicate a lot more with cloud servers, too, and infrastructure left around the city. In China, they are already building smart roads to get ready. Sensors along the side of the road watch for accidents or other problems and report to everyone what’s going on.

Speed limits can be changed, lanes opened or closed, and traffic light behavior changed by these smart roads, which will communicate with cars via 5G so cars can get real-time data about potholes, objects on the road, and changing traffic conditions ahead. Tesla is already starting to work like that, even without the smart road, by the way, because Tesla is making its own maps and updating those maps from its hundreds of thousands of cars. Many of these new “smart roads” will also have lots of cameras, for security, as well as to add new services. Your glasses and self-driving car might know exactly how many spaces are open on any given street before navigating you there, or where some of those electric scooters are so you will be able to continue your journey to your meetings for the day.

Interstates

As you drive along, you might notice that interstates between major cities are starting to look different. Hotels you might have seen in decades past are closing down, restaurants too, and they are moving to be near supercharging stations, which have changed traffic patterns through states like Kansas and Nevada. No longer do truckers need to stop to rest, so they’ve been flying by the hotels they used to stop at. They do need to charge their trucks and vehicles, though, so are stopping for 30 minutes at new charging stations, sometimes opting for a meal nearby, or some shopping for the kids. At least that will happen until the transportation industry totally automates and gets rid of all the humans using all those services today. That might start happening in the late 2020s, but regulators and policy-makers might hold humans into vehicles, mostly as a jobs act, until well into the 2030s.

Either way, we see countries like the United States getting a dramatic demographic shift over the next two to three decades due to all those interstate truckers being replaced by autonomous trucks and other changes, thanks to new “closer to home” warehouse farms and 3D printing plants, so lots of things may not need to be shipped such long distances in the future. 3D printing and Spatial Computing glasses fit together like peanut butter and jelly, too, since wearers can make and order 3D-printed objects.

Factories

Other, more traditional factories, like the Louisville Slugger bat factory in Louisville, probably won’t move from its historic location, but may see shipping costs dramatically reduce, and the time taken to get a product around the country might change due to autonomous vehicles so that they can open up new markets. The same will happen with a range of factories around America. Some companies, particularly drink or alcohol distributors, might be able to reduce the number of factories even further, and that could have an impact on communities.

Real Estate

You'll also see real estate developments further from city centers. Stockton, for instance, is about 100 miles from San Francisco. Now, residents there can take an electronic flying vehicle that looks like a drone (an eVTOL, as described earlier) for an affordable price directly into downtown San Francisco. The service can drop you off on top of one of a few buildings, and then you can take an Uber, Lyft, or maybe one of those new Zoox or Tesla robo-taxis to your final destination. Because this trip is less than 45 minutes, many people can see themselves commuting from Stockton to San Francisco; before flying drones, this would just be unfeasible. Basically, these are those big drones that can take off from small spaces, like a helicopter, albeit cheaper, quieter, and lighter, and will fly without a pilot. Sebastian Thrun says they will dramatically change the economics of moving people to cities that are around 100 miles away from city centers. Flying from Stockton to San Francisco, he told us, will take less than half an hour. In 2024, that trip can take two to four hours in a car, depending on whether you are stuck in traffic or not.

Remote Working

One thing we expect to see radical changes to is how often workers will be expected to come to headquarters to work. Today, a worker at, say, VMware in Palo Alto, which is right near Tesla and many other employers who do the same, might skip many commutes altogether, due to new Spatial Computing glasses that will let that employee collaborate with other workers almost as effectively as if those workers were in a conference room together. This will be one of the ways that employers can do something real about housing prices, as well as making real inroads into climate change-mandated carbon reductions.

Travel Times and Ranges

Other changes are predicted to occur where people stop along interstates. In a gas-powered manually driven car, humans can only drive about four hours before they start getting dangerously tired. Today, a trip from San Francisco to Reno is pretty safe for a single driver to do in one shot. As autonomous vehicles get better, though, particularly at interstate driving, you will be able to go far further as a single driver will not worry about falling asleep on a boring straight road. So, you'll see many people travel from San Francisco to, say, Wyoming, which takes about 20 hours of travel time. Or, Salt Lake City in 13 hours. Even including bathroom breaks and charging stops, you'll see America (and many places, like Canada and Europe, which have an electric charging infrastructure and mostly straight, well-painted roads) turn into a new kind of amusement park. "Hey honey, why don't we go skiing tomorrow in Salt Lake City?" could mean hopping into an autonomous vehicle at 6 p.m. for a 9 a.m. arrival at the lifts.

Navigation

Back to cities, though. Street cleaning and other infrastructure will use 3D maps of everything to navigate and also figure out objects, cars, and pedestrians to be avoided. Museums and other popular tourist destinations will have intricate AR audio tours. Already, Rome has audio tours from Rick Steves that you can listen to on your AirPods, and we used one tour in Yellowstone that automatically played as we drove around that National Park. It was magic and that was audio-only in 2018. Everything you look at will have an AR nameplate with details such as, “Here is the Grand Prismatic Spring—as you look around, we’ll show you the names of the mountains surrounding it and all the smaller geysers around too.”

Green Thinking

Much more cities will electrify in the 2020s, too, due to climate change concerns. China currently owns about 90 percent of the electric trains and buses in the world. We expect that mix to change as other countries start investing in electronic alternatives in their mass transit plans.

This climate change awareness will be the excuse many cities will use to electrify and automate all sorts of services, from mail to garbage pickup. After all, turning those things from diesel trucks with humans that aren’t as efficient to electric vehicles that are more efficiently operated by computers (and more safely, since they never will show up to work stoned or drunk) will be attractive to many cities as they struggle to meet carbon reduction goals and add new services for citizens wearing Spatial Computing glasses.

AI Assistants

Even as people use cars, they will use these glasses to join meetings, take classes, watch videos, play games, and get information on the world flowing by as their cars or drones take them places. Speaking of which, you will talk more and more to many computers almost at the same time. Where today you might say “Hey Siri” or “Hey Alexa,” sometime in the 2020s, you might start talking to just one of them, or a generic “Hey ‘cloud’” to command all services to figure out how to serve your query. Either way, by the end of the 2030s, we’ll be very used to speaking with computers, if we even think of them in that way anymore. Why do we say that? By 2030, AI systems will be producing many things for us, including music, food, and running our transportation system, and they also will be presented to us visually as new virtual, or “synthetic,” beings. Magic Leap first showed us its virtual being in 2019 at the **Games Developer Conference (GDC)** in San Francisco. Of course, that was heavily demoware, and couldn’t talk with users, or do all that much except play a simple game on a simple table.

We are sure that, by 2030, these virtual beings will actually sit in your car with you and will entertain you and assist you on your journey.

Entertainment

Speaking of entertainment, other changes in cities will be not just how cool they look. In late 2019, China had hundreds of drones flying over Guangzhou building a spectacular light show over its skyscrapers. That will seem quaint by 2030 as drones and AR are mixed into nightly entertainment on top of cities.

Raves, virtual wars, and more will play out every evening in city centers this way, while participants play new kinds of virtual games, which will make the 2020 version of *Pokémon GO* or Minecraft Earth seem pretty simple indeed. How? This will be aided by all the data that companies are grabbing off of cars, cities themselves, and those AR glasses. There's more on Spatial Computing in entertainment in our next chapter, *Chapter 5*.

Fewer Cars

Finally, if you are walking around a modern city, you might notice two things: no more parking spots and no more cars in a city center, with a preference for bikes, electric scooters, or other public transportation, maybe including robo-taxis. Why? Europe is showing the way—as you get rid of cars, you make cities much more livable and attractive, not to mention safer. The city of the future will help you live better and will be designed by data that proves it so.

“Data Bubbles” Will Abound

Put it all together and we see the cities of the future with millions of newly-equipped autonomous vehicles on the road within a few years. In the San Francisco Bay Area, it already is very common to have two or three Teslas around you, with a Cruise from General Motors or a Waymo driving nearby, too. These are all like little “data bubbles.” Each can see a couple of hundred meters around the car, building a 3D model, and their AI systems convert all that into component parts, like lane markings, stop signs, other cars, and pedestrians. These data bubbles soon will enable very detailed new maps, and bring radical changes to our cities as parking spaces get converted into drop-off zones, among other changes. Not to mention that we'll be wearing Spatial Computing glasses to summon new services from rides to food delivery because of the autonomous vehicles on the road.

This will bring deep shifts in our culture. Back in the 1960s, we sang “we drove our Chevy to the levy and the levy was dry.” By 2030, we’ll need a whole new set of lyrics. What a time to be alive!

In our next chapter, we’ll attack the **Technology, Media, and Telecommunications (TMT)** vertical—one that goes to the heart of the Spatial Computing industry.

5

Vision Two: Virtual Worlds Appear

As we did in the previous chapter in *Part II*, we will be looking at an industry vertical in this chapter. This chapter's vertical is **Technology, Media, and Telecommunications (TMT)**. Since this vertical goes to the core of what Spatial Computing is very much about, it will necessarily be longer than the other chapters that focus on industry verticals.

New realities are appearing thanks to **Virtual Reality (VR)** and **Augmented Reality (AR)**. In a few years, though, the capabilities of VR and AR will change and morph into more sophisticated Spatial Computing glasses, which will arrive with massive increases in bandwidth thanks to 5G and new **Artificial Intelligence (AI)** capabilities, which are due in part to the R&D being done on autonomous vehicles. Here, we dive into some of the fundamentals of the TMT vertical, as well as the profound changes we expect.

From 2D to 3D

As we look back at the technology industry, especially when it comes to personal computers and mobile phones, which customarily present 2D images on a flat screen, you can see that there's an underlying set of goals—beliefs, even. These include things such as connecting everyone together, giving you “superpowers” to analyze business data with spreadsheets, and the ability to communicate better with photos and videos. Previous generations were just about giving us better tools than the mechanical ones that existed before, whether the printing press or the old rotary phones that so many of us now-old people had in our homes.

Increasing the Bandwidth

Elon Musk puts it best: the goal is to increase the bandwidth to and from each human. Now, he's investing in a new kind of computing with his Neuralink company, that includes "jacking in" and hooking a computer directly up to your brain. That will take a lot longer to show up for most people than Spatial Computing, simply because of the cost of opening up a human skull and having a robot surgically implant tiny wires directly in a brain.

One way to look at it is via our current computing interfaces. If you have a thought that you want to communicate with other people, you probably use a mouse or a finger on a screen to open an application or a window, and you use a keyboard to type a message. That's pretty much the same whether you want to communicate in a Microsoft Word document, an SMS text message, an email, or a post on X (previously Twitter). On the other side, your reader sees what you wrote on a screen and that person's eyes and visual perception system translate that into something their brain can understand, process, store, integrate, and then maybe reply to.

That process of reading is pretty fast in comparison to the process of writing. There's friction in communication. We're feeling it right now as we type these words into Microsoft Word and are trying to imagine you reading this book.

The process isn't a whole lot faster even if we use video. After all, you can even have someone, or better yet, AI system, read this text to you while you walk down the street. That would be hard to do if we made a video. (Note that we previously discussed the relevance of AI, including machine learning, deep learning, and computer vision, in *Chapter 2, Four Paradigms and Six Technologies*, so we won't replicate that information here.)

Plus, you can skim and do other things with text that video makes very hard. Your mind, on the other hand, would far rather watch a movie than read a movie script, or watch a football game on TV than read a report about it later.

Where we are going with this is that the tech industry is spending billions of dollars in an attempt to make it easier to communicate and work with each other, or even to better enjoy the natural world. We see the ultimate effect of this in the "jacking-in" fantasy that Elon Musk and others are trying to make a reality. Some take a more negative view, like John von Neumann and his term "the singularity," which was later popularized by Vernor Vinge in his 1993 essay in which he wrote that it would signal the end of the human era. Scary stuff, but the need to communicate our thoughts with not just other people, but machines, like we are doing now with Siri, Alexa, and Google Assistant, is opening up new ways of communicating with machines.

Elon Musk's Neuralink, which uses a robot to hook up tiny wires directly to a human brain, is the most forward approach we've seen actually being put to use. That said, we don't think we'll sign up for that in the early exploratory years of the company!

In its current form, such technology brings with it deep side effects and, as of today, costs hundreds of thousands of dollars due to the necessity of a surgical procedure to install something like Neuralink. If you have, for example, Parkinson's, you may get such a technology implanted by the end of the 2020s, but the rest of us will be stuck wearing some form of on-the-face computing, whether eyeglasses or contact lenses, for the foreseeable future. It will take longer than the end of the 2020s to get full FDA approval for Neuralink outside of particular medical conditions having earlier approval.

Change is Coming

Spatial Computing is evolving quickly, with several different device types (which we'll go into more depth about in a bit), but they all require putting screens on your face. We dream of a lightweight pair of glasses that will bring the whole spectrum of Spatial Computing experiences that cover the reality-virtuality continuum introduced by researcher Paul Milgram, when he said that soon humans would be able to experience extended reality, or XR, along a spectrum from "completely real" on one side to "completely virtual" on the other.

Soon, we'll see devices that let you switch between real and virtual so quickly and seamlessly that some might start to lose track of the reality of what they are experiencing. You could say we have that already with Apple Vision Pro. We'll cover the pros and cons of that in later chapters but it's also good for us to slow down and revisit the technology that's here today so that we can have a decent conversation about where things have been and where they might be going.

We argue that it's easier for you to understand what, for example, a Macintosh could do if you knew what previous computers could do and the design constraints that were placed on earlier devices and ecosystems. Our kids forget why the Dodge tool in Adobe Photoshop looks like a lollipop, but those of us who used to work in darkrooms with chemicals and enlargers remember using a tool like that to make parts of photos lighter. The same will happen here.

First, we constantly get pushback every time we explain that you will soon be forced to wear something on your face to fully appreciate the higher communication speed that Spatial Computing brings to your brain.

Yes, why glasses? Why can't we get to this "high bandwidth to and from the brain" world with standard monitors?

Thomas Furness, the American inventor and VR pioneer, asked himself just that question back in the 1960s when he was designing cockpits for pilots in the military. He, and many others, have tried over the years to come up with better monitors. The US government funded HDTV research for just that reason. The answer comes from how we evolved over millions of years. Evolution, or, some might say, God, gave us two eyes that perceive the analog world in three dimensions, and a large part of our brains are dedicated to processing the information our eyes take in.

If we want to take in more and get the bandwidth to our brains higher, we need to move our monitors onto our eyes. If we want to improve the outbound bandwidth from our brains to the computing world, we need to use voice, hands, eyes, and movement and add more sensors to let us communicate with computers with just our thoughts. All of this requires wearing sensors or microphones either on our faces or heads. There is no way around this except in special situations. Andy Wilson, one of Microsoft's top human/machine researchers up in Redmond, Washington, has a lab with dozens of projectors and computers where he can do a lot of this kind of stuff in one room—sort of like the Holodeck in Star Trek. But even if his room is successful, most of us don't want to compute only while sitting in a singular place, and most humans simply don't have the space in their homes and offices for such a contraption anyway.

We see this in high-end uses already. The pilots who fly the F-35 fighter jets tell us they are forced to wear an AR helmet simply to fly the plane. "I'll never lose to an F-16," one pilot at Nellis Airforce Base told us.

"Why not?" we asked.

"Because I can see you, thanks to my radar systems and the fact that your plane isn't stealthy, but you can't see me. Oh, and I can stop and you can't."

Soon, many of us will be forced to follow the F-35 pilot in order to maximize our experiences in life. Those of us who already wear corrective lenses, which is about 60 percent of us, will have an easier time since we already are forced to wear glasses simply to see. We see a world coming, though, where Spatial Computing devices bring so much utility that everyone will wear them, at least occasionally.

This will bring profound changes to humanity. We are already seeing that in VR. Users of VR are already experiencing a new form of entertainment and media; one that promises to be a bigger wave than, say, the invention of cinema, TV, or radio in previous generations.

This new wave builds on the shoulders of those previous innovations and even subsumes them all: wearing the glasses of 2027, you will still be able to read a magazine, watch a TV show, or listen to your favorite radio station (or the podcasts that replace them). As all of these improve, we are seeing new kinds of experiences, new kinds of services, and even a new culture evolving.

To those whom we have told this story, much of this seems improbable, weird, dystopian, and scary. We have heard that with each of the previous paradigm shifts too, but because this puts computing on your face, the resistance is a bit more emotional. We see a world coming over the next few years where we mostly will get over it.

Emergence as a Human Hook

So, why do you need to put on a headset? Well, when you're inside the headset, you'll be able to experience things you've never been able to experience on a computer monitor before. But what's the purpose?

The Brain and VR

Why do human beings like playing in VR? Emergence is one reason, and we'll cover a couple of others in a bit. What is emergence? When applied to video games, emergence is the feeling you get from playing complex situations, and unexpected interactions show up from relatively simple game dynamics. The same feelings are why we go to movies. It's why we like taking regular drives on a curvy road. It's why we like holding a baby for the first time. It's why we like falling in love. All of these create positive chemical reactions in our brains that are similar to the emergence that happens in video games. VR enables this more frequently. Since the 1970s and 1980s, board games and table-top role-playing games, such as *Cosmic Encounter* or *Dungeons and Dragons*, have featured intentional emergence as a primary game function by supplying players with relatively simple rules or frameworks for play that intentionally encouraged them to explore creative strategies or interactions and exploit them toward the achievement of victory or a given goal. VR makes it that much more powerful.

If we have sharper and more capable screens, we can experience things that make our brains happier, or, if you will, a flow state, which is the euphoria that both workers and surfers report as they do something that brings high amounts of emergence.

The Right Tool for the Job

The thing is, we're going to need different kinds of headsets for different things. VR headsets separate us from others. They aren't appropriate to wear in a shopping mall, on a date night, or in a business meeting. But, in each of these situations, we still want digital displays to make our lives better and to give us that feeling of emergence.

It's the same reason we would rather watch a football game on a high-resolution HDTV, instead of the low-resolution visuals of our grandparents' TVs. It's why we tend to upgrade our phones every few years to get better screens and faster processors. When we watch sports or movies on higher-resolution screens, bigger screens, or both, we can experience emergence at a higher rate.

Emergence is also why notifications on our phones are so darn addictive. Every time a new notification arrives, a new hit of dopamine arrives with it. Marketers and social networks have used these game dynamics to get us addicted and build massive businesses by keeping us staring at our feeds of new items. VR promises to take our dopamine levels from these mechanisms to new levels. In later chapters, we'll cover some of the potential downsides of doing that, but here, we'll discover why Spatial Computing devices are seen as a more powerful way to cause these dopamine hits, and thus, a promise to build massive businesses.

These dopamine hits, along with powerful new ways to be productive, are coming from a range of new devices—everything from tiny little screens that bring new forms of your phone's notification streams, to devices that so completely fool your brain that they are being used to train airplane pilots, police, and retail workers.

The “Glassholes” Show the Way

It started with a jump out of a dirigible at a Google programming conference back in May 2014. The jumpers were wearing a new kind of computer: one that put a little screen and a camera right near their right eye. They broadcast a video feed from them to everyone. People jumped up to be the first to order one (we were amongst the first in line) because the demo was so compelling.¹

Defeated Expectations

That said, the demo oversold what would actually materialize. Before we got ours, we thought it would be an experience so futuristic that we had to put down \$1,500 to be the first to experience it. Now, don't get us wrong, the first year of wearing Google Glass was pretty fun, mostly. People were highly interested in seeing it. At the NextWeb conference, attendees stood in line for an hour to try ours out.

Most people walked away thinking they had seen the future, even though all they saw was a tiny screen that barely worked when we talked to it with “Hey, Google” commands.



Figure 5.1: Photo credit: Maryam Scoble. This photo of coauthor Robert Scoble wearing Google Glass in the shower went viral and has been on BBC TV many times, along with many other media outlets.

As we got to wear ours, it became clear that Google Glass wasn’t that futuristic and certainly didn’t meet expectations. Robert Scoble’s wife, Maryam, upon seeing it for the first time, asked, “Will it show me anything about my friends?” The truth is: no, it didn’t, but that was an expectation. There were other expectations too, set up by that first demo. The camera wasn’t that good. Certainly, it wasn’t as good as everyone thought it was, which later became a problem since people thought it was constantly recording and uploading, or, worse, doing facial recognition on everyone.

The expectations were, though, that it was a magical device that you could look at anything with and learn more about that person or thing. Those expectations were never really corrected by either Google or the people who owned them, who soon gained a derisive name: “Glassholes.” This derisive tone came about after a journalist named Sarah Slocum wore a pair into a bar and was, in her words, accosted due to wearing them. After that, we saw more businesses putting up signs forbidding users from wearing them, and Google soon gave up trying to sell them to consumers.

Where Did Things Go Wrong?

We are frequently asked, how did Google Glass go so wrong? Also, what does it mean for the next wave of products about to arrive?



Figure 5.2: Photo credit: Google. Here is a closer look at the original Google Glass, which had a small transparent screen, along with a camera, microphone, speaker, and battery.

We think it mostly came down to that dramatic launch and the expectations it caused, but there was a basket of new problems that humans had never considered until Google forced this into public view. Doing such a spectacular launch with people jumping out of dirigibles over a San Francisco convention center with thousands of people in the audience, all watching on huge screens, with millions of other viewers around the world, just set a bunch of expectations that it couldn't meet.

We list the main complaints many had here:

- The screen was too small.
- The 3D sensors weren't yet ready.
- The **Graphics Processing Units (GPUs)** were very underpowered so the graphics that it could show were very rudimentary.
- The battery life wasn't great.
- The price of \$1,500 was relatively expensive for the features it provided.
- It didn't have screens or cameras for both eyes, so it wasn't appropriate for doing more serious AR like what Microsoft's HoloLens 2 does today.
- The cameras weren't as good as in your standard GoPro.

- The screens were of far lower resolution and size than even your cheapest phone at the time (and compared to today's devices, are like a postage stamp compared to a letter-sized piece of paper).

Worse yet, Google didn't expect the social contract problems that showed up months after they shipped. Humans had evolved to look into each other's eyes to verify and communicate trust, and there was a new device messing with that. Many people focused negative attention on the camera. This touched on privacy expectations that hadn't been explored by previous technologies, and the lack of utility or affordability just added to a perfect storm of negative PR that it couldn't escape from.

That said, Google Glass was, and is, an important product that set the table for a whole range of new products. It was the first computer that you wore on your face that most people heard about, at a minimum, and got a whole industry thinking about putting computers on people's faces. But, this is a hard space to get right. Why? Because it's so damn personal. The computer is on your face in front of your eyes. This isn't like a desktop computer that you can hide in an office and only use when you aren't trying to be social. When fully expressed, these new computing forms change what it means to compute and what it means to be human.

What It Means to Be Human

Whoa. Change what it means to be human? Yes. Soon, you'll be able to do things with these glasses that futurists have been predicting for decades. This includes everything from ubiquitous computing to next-generation AR that changes how we perceive the real world in deep ways. Qualcomm's new prototypes have seven cameras on a device about the size of or smaller than Google Glass, which only had one camera. One of these cameras watches your mouth with AI that was barely being conceived when Google Glass was announced. The others watch your eyes and the real world in high detail that bears little resemblance to what Glass' low-resolution camera was able to do.

It's important for us to pause a moment and take a bird's-eye view of the Spatial Computing landscape though, and get a good look at all the devices and where they came from. Spatial Computing is a big tent that includes products that are small and lightweight, like Google Glass, all the way to devices like the Varjo, which, in 2020, opened with an initial cost of \$10,000 and needed to be tethered to a \$3,000 PC. It now has a working version, the Varjo VR3, for enterprise use that costs \$2,500. The Varjo is big and heavy, but the experience of wearing one is totally futuristic compared to the early Google Glass prototypes.

No longer are we stuck with a tiny postage stamp of a screen. Inside the Varjo, you have huge high-resolution screens. Where Google Glass is like a mobile TV on a kitchen counter, Varjo's device is like being in an IMAX theater. With it, you can simulate flying in an Airbus Jet's cockpit. Where the Google Glass could barely show a few words of text, the Varjo lets you feel like you are in the cockpit as a pilot, with tons of displays blinking away and everything so photorealistic that you feel like you are looking at hundreds of real knobs and levers in a real plane. Out the window, while wearing the Varjo (which would be the best experience for a headset), you can see cities and landscapes emerge, along with runways from a distant airport you were supposed to land at. All simulated, yes, but your brain sees so much resolution it thinks it is real.



Figure 5.3: Photo credit: Varjo. The highest-resolution VR headset, as of April 2024, is the Varjo XR-4.

The Focal Edition version, which is priced at \$9,990, is used by Volvo to design and test future car designs.

As the 2020s evolve, you'll see a spectrum of devices with far more features and capabilities soon escape from R&D labs—including a small pair of glasses. The Ray-Ban Meta Wayfarer is a good example and selling in numbers much more than earlier expected, and one that's similar to what Google Glass tried to do. On the other side of the spectrum are VR headsets from Varjo, Meta, and others. We'll detail them here to give you context on the market and the kinds of devices that fit under the Spatial Computing tent.

When we say these will change what it means to be human, we mean that soon, the way that you experience the real world will be completely changed by AR. Already, Snapchat is changing faces, floors, and buildings with its version of AR, which is technically no AR at all.

Soon, every part of your human experience will be augmented as new Spatial Computing glasses, headsets, and other devices appear. The Varjo is giving us an early taste of just how complete these changes could be—inside the headset, you experience a virtual world in such detail that you no longer want to turn it off. It is being used to train pilots, and when we got a demo, the virtual cockpit that was presented was so sharp and colorful that we imagined that someday we'd just experience the whole world that way. These changes will deeply change human life. We predict that the changes will be deeper than those brought to us by the automobile or telephone.

Embodiment, Immersion, Presence, Empathy, and Emergence

In February 2014, Mark Zuckerberg visited Stanford's Virtual Human Interaction Lab. Prof. Jeremy Bailenson and his students walked Facebook's founder and CEO through a variety of demos of how VR could change people's experiences. He walked across a plank that made you feel fear of falling off. He experienced how VR could help with a variety of brain ailments and played a few games. Bailenson ran us through the same demos years later and said we had the same experiences, including being freaked out by finding ourselves standing on a narrow piece of wood after the floor virtually dropped away to reveal a deep gap. Bailenson adeptly demonstrates how VR enables powerful emergence, along with a few new ways to cause our brains to be entertained: embodiment and immersion, not to mention a few other emotions, like the vertigo felt on that wooden plank.

Embodiment and Immersion

Embodiment means you can take the form of either another person or an animal. Bailenson showed us how embodiment could give us deeper empathy for others. In one demonstration, we experienced life as a black child. In another demo, we had the experience of being a homeless man who was harassed on a bus. Chris Milk, a VR developer, turned us into a variety of forms in his experience, "Life of Us," which he premiered at the Sundance Film Festival in 2017. We watched as people instantly flew as they discovered that they had wings, among other forms, from being tadpoles to stock traders running down Wall Street.

Immersion is the feeling that you are experiencing something real, but in a totally virtual world, in other words, being present in that world. The perception is created by surrounding the user of the VR system with images, sound, and other stimuli that provide an engrossing total environment.



Figure 5.4: Photo credit: Facebook. The \$400 Oculus (now Meta) Quest changed the market by making VR far easier thanks to having no cords and no need for a powerful PC to play high-end VR games and other experiences.

Within a few months of getting the demos from Bailenson and his team, Zuckerberg would acquire Oculus for \$2 billion. Oculus was started a couple of years prior in Irvine, California, by Palmer Luckey, Brendan Iribe, Michael Antonov, and Nate Mitchell. The early VR device was a headset that had specialized displays, positional audio, and an infrared tracking system. Sold on Kickstarter, Oculus raised \$2.4 million, which was 10 times what was expected.

Evolution

Since 2012, Oculus has evolved and now is called Meta. The headset was paired with controllers that let a user touch, grab, poke, move, shoot, slap, and do a few other things inside virtual worlds. A new term, 6DOF, got popular around the same time. That stands for Six Degrees of Freedom and means that the headset and controllers could completely move inside a virtual world. Previous headsets, including the Oculus Go, were 3DOF, which meant that while you could turn your head, you couldn't physically move left, right, forward, or backward. To experience the full range of things in VR, you'll want a true 6DOF headset and controllers.

Systems that aren't 6DOF won't deliver the full benefits of immersion, presence, embodiment, emergence, or the empathy that they can bring. That said, there are quite a few cheaper headsets on the market and they are plenty capable of a range of tasks. 3DOF headsets were great for viewing 360-degree videos, or lightweight VR for educational uses that don't require navigating around a virtual world, or the other benefits that 6DOF brings. It is this addition of 6DOF to both the controllers and the headset that really set the stage for the paradigm shift that Spatial Computing will deliver.



Figure 5.5: Photo credit: Robert Scoble. Meta had huge signs when Oculus Quest first showed up, showing the market advantages Meta has in getting its products featured.

The product that got the most accolades in 2019 was the \$400 Oculus Quest from Facebook (now Meta). This opened up VR to lots of new audiences, before it came out most of the other VR headsets, and certainly the 6DOF ones, were “tethered” to a powerful PC. Our first VR system was an HTC VIVE connected to a PC with an Nvidia GPU, which cost more than \$2,500. Today, for most uses, a \$500 Quest 3 is just as good, and for the few more serious video games or design apps that need more graphic power, you can still tether to a more powerful PC with an optical cable that costs about \$80.

Meta Has the Edge

Meta has other unnatural advantages, namely its two billion users, along with the social graph it brings (Meta has a contact list of all of those users and their friends). That gives Facebook the market power to subsidize VR devices for years and keep their prices down. The closest market competitor, Sony’s PlayStation VR, is a more expensive proposition and requires a tether to the console.

The Quest comes with a different business model, though, that some find concerning: advertising. In late 2019, Facebook announced it would use data collected on its VR headsets to improve the advertising that its various services show its users.



Figure 5.6: Photo credit: Robert Scoble. Here, Scoble's son, Ryan, plays on Oculus Quest while his dad charges his Tesla at a supercharger.

This gives Meta a near monopoly in sub-\$500 devices that bring full-blown VR and it is the only one that brings this in a non-tethered way.

What is notable is the shifting of price points. When PlayStation VR first came on the market, it was a lot cheaper for a kid to get when compared to buying a \$2,000 gaming PC and a \$900 headset, like the HTC VIVE. Anyone who wants to go mainstream has to figure out how to deal with that price differential due to the advertising-subsidized pricing that Meta has. Google could also do the same, but its forays into Spatial Computing, including Glass and the Gear VR headset systems that used mobile phones, fell flat with consumers because they didn't provide the full VR magic of embodiment, immersion, presence, and so on.

Market Evolution

Our prediction is that Meta will consolidate much of the VR market due to its pricing, social element, and the advantages that come with having a war chest of billions of dollars to put into content. Its purchase of Beat Saber's owner, Beat Games, in late 2019 showed that it was willing to buy companies that make the most popular games on its platform, at least in part to keep those games off of other platforms that will soon show up. Beat Saber was the top-selling game in 2019, one where you slice cubes flying at you with lightsabers. This consolidation won't hit all pieces of the market, though. Some will resist getting into Meta's world due to its continued need for your data to be able to target ads to you. Others, including enterprise companies, will want to stay away from advertising-supported platforms. For that, we visited VMware's Spatial Computing lab in its Palo Alto headquarters.



Figure 5.7: Photo credit: Robert Scoble. A wide range of headsets hang on the wall inside VMware's Spatial Computing lab in its Palo Alto headquarters.

At VMware, they buy and test all the headsets, so they can see all the different approaches various VR and AR headset manufacturers are taking. VMware's main business is selling virtual machines to run the world's internet infrastructure inside data centers and manage all the personal computers that connect to it.

Why is it investing in Spatial Computing? Because it is building an array of management infrastructure for enterprises to use and sees a rapidly growing demand for Spatial Computing. Walmart, for instance, bought more than 10,000 VR headsets for its training initiatives. Simply loading an **operating system (OS)** on that many headsets is a daunting challenge—even more so if corporate IT managers want to keep them up to date, keep them focused on one experience, and make sure that they haven't been hacked and have access to appropriate things on corporate networks and other workers.

Inside VMware Labs, Matt Coppinger, the director, walked us through the headsets and the work that they are doing with both VR and AR. He told us about enterprises who decided against going with Facebook's solution because Facebook wanted to control the software loaded on the headset too tightly. Lots of enterprises want to control everything for a variety of reasons, from the ease of use that comes from getting rid of everything but one app that's already running when you put on the headset, to being able to control which corporations the collected personal data is shared with.

He split up the market into different use cases (3DOF not included):

- **6DOF tethered/high end:** This means the Varjo XR-4, with its ultra-high-resolution and low-latency reign, which is why Volvo uses it for car design. The headsets at the high end have sharper screens and a wider field of view, which makes them the best choice for architecture, design, and use cases where you have highly detailed and complex things to work on, like an Airbus cockpit simulation for teaching pilots to fly.
- **6DOF tethered/mid-market and lower market:** The Valve Index and HTC VIVE headset products fit here. These tethered headsets allow the use of Nvidia's latest GPUs, so are great for working on 3D design, architecture, and factory work but are much more affordable than the \$10,000 Varjo. Most of these will cost you \$1,500 to \$3,000 for a decent PC and about \$1,000 for the headset systems.
- **6DOF self-contained:** The Meta Quest 3. This is the most exciting part of the VR market, they say, because it brings true VR at an affordable price with the least amount of hassle due to no cord and no external PC being required.

Keep in mind we have been through many waves of VR devices. VR has existed in the military since the mid-1960s. In the 1990s, people had VR working on Silicon Graphics machines (if you know what those are, that dates you to at least 40 years old or so). We flew plane simulators on such systems. But the VR form factor back then was huge and a Silicon Graphics machine cost more than \$100,000.

We tend to think of that phase as an R&D phase, where the only uses possible were either military or experiences that cost a lot of money per minute. One of our friends opened a retail shop where you could pay \$50 for a few minutes inside one to simulate flying a fighter jet. Since most of us couldn't afford the fuel for a real jet, this was as close as we were going to get, and it was pretty fun. Today's version of the same is a **Location-Based Entertainment (LBE)** experience, where you can pay \$20 to \$50 for an average of 15 minutes per person to be "inside a movie" at Sandbox VR and others.

2014's New Wave

Instead of looking so far back, though, we see the new wave of Spatial Computing really started in 2014 when Facebook bought Oculus. That purchase by Mark Zuckerberg really started up a whole set of new projects and caused the world to pay attention to VR with fresh eyes.

We still remember getting a “Crescent Bay” demo, one of the first demos available on the Oculus Rift headset, in the back room at the Web Summit Conference in Ireland in 2016. That headset was tethered to a PC. Two little sensors were on a table in front of us. Once we were in the headset, they handed us controllers. It started out in darkness. When the demo started, we found ourselves on the top of a skyscraper.

It was the first time we had felt vertigo; our brains freaked out because it was so real-feeling due to the immersive power of VR that we thought we might fall to our deaths.



Figure 5.8: Photo credit: Meta. The original Oculus Rift VR system, with its two sensors and controllers.

Not seen here is the PC with a big GPU inside, usually made by Nvidia, that it's tethered to.

Now, how did a computer make us scared of falling to our death? Well, the headset, as it was moved around, showed a 360-degree world. We felt like we were really on top of a skyscraper and could fall over the edge. This was far different than any computer or medium we had ever seen or experienced before. The power of putting highly tracked screens on our faces that could reveal a virtual world was like magic. We barely noticed the cord that was hanging from a device over our head leading to a PC on the side of our play area.

Later, after we bought our own, we would get the limitations of the sensor system. It built a virtual room that we couldn't really move out of. In our case, in our homes, it was a virtual box that was a few feet wide by a few feet long. The cord and the sensors kept us from leaving that box and those limitations started to chafe. We couldn't play VR while on the road traveling. We couldn't take it to shopping malls. Trying to show people VR in their backyards proved very difficult.

Getting even two of the systems to work meant getting black curtains to separate the two systems since the sensors didn't like being next to another set that was trying to operate.



Figure 5.9: Photo credit: HTC. The HTC VIVE and its Lighthouses.

The HTC VIVE was popular because of its superior tracking system. These black boxes, called “Lighthouses,” were aimed at a VR player and sprayed a pattern of invisible light that let the system track the headset and controllers accurately. Let’s talk about those sensors because how tracking is done is a key differentiator between headsets. HTC VIVE is still a favorite amongst heavy users of VR, in part, because its tracking system is so good. It consists of two black boxes (you could use more if needed, in, say, a large space) that spray invisible light on a VR user walking around. This invisible light is used to track headsets, controllers, and newer trackers that users can put on objects, their feet, and other things. These black boxes, called “Base Stations,” are better at tracking than the ones Meta sold with its Oculus Rift system. The original Oculus ones were simpler cameras that weren’t quite as exact and couldn’t track as many things as the HTC system did.



Figure 5.10: Photo credit: Robert Scoble. A girl plays with an HTC VIVE. Here, you see the tether to a gaming PC.

A lot of these advantages, though, have been whittled away at by the Oculus team, who shipped quite a few updates to make its system work better and, finally with the Oculus Quest, this older “outside-in” tracking system was replaced by a series of four cameras on the front that implemented the newer tracking system that didn’t require putting boxes or cameras around a VR user.

Inside-Out

We believe this newer system, which industry people call “inside-out” tracking because the sensors are looking from the headset “out” into the real world, is better for Spatial Computing because it frees users from having to calibrate sensors and place boxes or cameras around, and it also means you have the freedom to walk around much bigger spaces, leading to wearing glasses in a few years where you can play virtually anywhere without worrying about the tracking.

The Oculus “inside-out” tracking had also shown up in its Rift S model, which replaces the earlier Rift, with its two external trackers. What competed with the Rift S was the Valve Index, which got many gamers to salivate with its higher specs, including a sharp set of 2880x1600 monitors with a 130-degree field of view. The Oculus Quest, for comparison, only had 1440x1200 pixels per eye, a slower 72 Hz refresh rate, and a smaller field of view at around 100 degrees. The negative for all that performance is that the Valve requires a user to be tethered to an expensive gaming PC.



Figure 5.11: Photo credit: Valve. The Valve Index got a lot of great reviews from gamers, who valued its high-resolution screens with a wide field of view.

Finally, the road to bringing us this magic has been quite rocky. Two notable failures are Google Daydream and Microsoft's Mixed Reality licensed headsets.

Why did these do so poorly in the market? Well, for different reasons.

Mistakes Were Made

The Google Daydream headset, and a similar idea from Samsung, let users plop their phones into a headset. This idea promised low-cost VR, but it didn't deliver very well. First of all, the mobile phones that were on the market three years ago didn't have hugely powerful GPUs, so the graphic capabilities were muted. Worse, mobile phones often got too hot because of this underpowered nature, and VR pushed the phone's capabilities harder than anything else could, draining the batteries quickly on devices people needed for a lot more than just playing in VR.

Ours crashed a lot, too, due to the processing demands that heated everything up. The weight of the phone, too, meant that these headsets weren't very comfortable, with a lot of weight out in front of your nose.

What really killed them, though, was that they were only 3DOF, there just wasn't enough content to view, and the people who might have hyped them up had just gotten their full 6DOF systems from HTC or Facebook and weren't likely to spend hours playing underpowered mobile-based headsets.



Figure 5.12: Photo credit: Google. One of the attempts that didn't go well with consumers was the Google Daydream headset, which let you drop your mobile phone into this device to do VR. But it wasn't very full-featured; it couldn't do 6DOF, for instance, and only ever had very limited content developed for it.

The Microsoft Mixed Reality headsets tried to fix many of these problems but also failed due to a lack of customers, mostly because the people who really cared about VR saw these attempts as underpowered, with poor sensors, and not having much marketing behind them. Translation: the cool games and experiences were showing up on the HTC VIVE or Oculus Rifts. Things like Job Simulator, which every child under the age of 13 knew about due to the brilliant influencer marketing of its owner, Owlchemy Labs, which now is owned by Google, were selling like hot-cakes, and games like that needed the more powerful headsets.

We tend to think it was a mistake, too, for Microsoft to call these VR headsets “Mixed Reality.” That confused the market since they came after Microsoft had released the critically acclaimed HoloLens, which we’ll talk about a lot more in a bit, which was a true Mixed Reality experience where monsters could come out of your actual walls. These VR devices weren’t that, confused the market, and made the market wonder why they weren’t just saying these are VR headsets. Either way, these are gone now except for the Samsung HMD Odyssey, and the Oculus Quest has set the scene for the next wave of Spatial Computing products.

We are seeing products in R&D that bear very little resemblance to the big, black headsets that are on the market as the 2020s open up. In fact, at CES, Panasonic showed off a new VR prototype that looked like a steampunk pair of lenses and was much smaller than the Oculus Rift. Oculus co-founder Palmer Luckey made some X excitement when he posted a photo and praise of the Panasonic (which we doubt will do well in the market due to its poor viewing angle, lack of content, and lack of 6DOF).

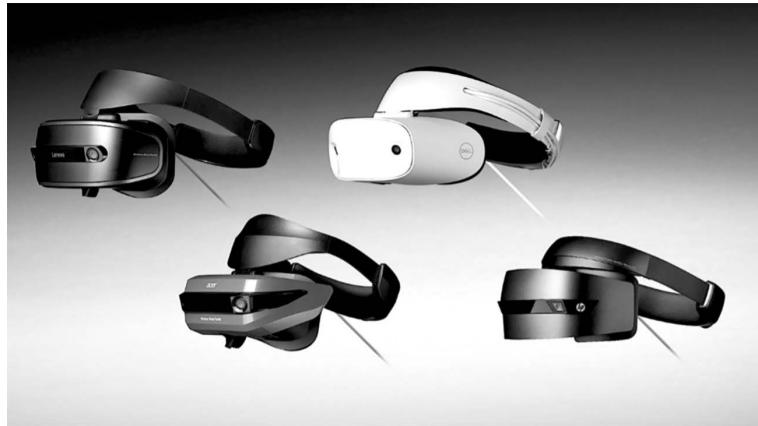


Figure 5.13: Photo credit: Microsoft. The original Microsoft-licensed VR headsets from Lenovo, Dell, Acer, and HP [from left to right and top to bottom]. The Samsung HMD Odyssey is missing from the photo, as it came out many months later than these, was relatively more expensive, and had the reputation of being the best quality of the bunch. Only the Odyssey is still available for direct purchase at a \$499 retail price.

Either way, VR is here now, and the magic it brings of immersion, embodiment, and presence sets the stage for a whole range of new capabilities that will show up over the next few years.

The Spectrum: Augmented Reality from Mobile to Mixed Reality

The popularity of Amazon's Alexa and its competitors shows us that humans are hungry for a new kind of computing. This form of computing lets us see new layers of information on top of the existing world. It is computing that surrounds us, walks with us through shopping malls, and assists us while working out. This kind of computing could potentially even save our lives as it collects our visual, biometric, and textual data.

Close Computers

To make it possible to get this level of human-machine interface, we will need to get computers closer to our brains and certainly our eyes.

We can do some of this kind of computing by holding a phone in our hands, which is already used for AR by kids playing *Pokémon GO*. Holding a mobile phone, though, won't get us to the promised land of true Spatial Computing. It must be said, however, that developers are using mobile technology to build a variety of things that are important to pay attention to.

As the 2020s open, we are seeing some of these wearables evolving in the marketplace with a spectrum of devices.

On one side, we have very lightweight devices that look similar to the Google Glass device. A good example of these lightweight devices is the Ray-Ban Meta Wayfarers, but there are a number of them from Vuzix, Epson, XREAL, and others. These have screens that show you a variety of information overlaid on top of the real world and do it with a form factor that looks like a pair of sunglasses.

At the other end of the spectrum are devices that completely change our visual experience as we move around in the real world. These devices, which include Microsoft's HoloLens 2, could be seen as a kind of "mixed-reality" device that brings advanced graphical features and a ton of sensors along with a new form and AI to see, process, and display all sorts of new experiences, both visual and auditory, that profoundly changes how we interact with the real world.

It is these higher-end devices that have us dreaming of a radically new way of computing, which is what we are covering in this book: Spatial Computing. With these devices, you're wearing seven cameras along with a bunch of other sensors and because of the amount of processing that needs to be done and the kinds of visual displays that need to be included, these devices tend to be both heavier and more expensive than devices, for instance, like the Ray-Ban Meta Wayfarers.

Let's go back to where AR came from.

The Origins and Limitations of Augmented Reality

For us, the potential for AR came up in 2011 as Peter Meier, then CTO of a small AR company in Munich, Germany, took us into the snow outside of its offices and showed us dragons on the building across the street. These virtual dragons were amazing, but the demo was done with a webcam on a laptop. This is hardly easy to do for consumers. Earlier attempts, like Layar, which was an AR platform for mobile phones, didn't hit with consumers because the mobile phones of the day were underpowered, and most didn't appreciate AR on small screens. These early efforts, though, did wake up a lot of people, including us, to the possibilities of AR. It also woke Apple up, who started working earnestly on Spatial Computing.

Within a few years, Apple had purchased that company, Metaio, and Peter today works there. Since then, Apple has purchased dozens of companies that do various pieces of Spatial Computing, and that acquisition trend continues into 2020. While we were typing this, Apple bought another company that does AI that will prove important to the future we describe here.

Metaio's work was the foundation of a new set of capabilities built into iPhones. Developers know this AR framework as "ARKit." Google has a competitive framework for Android users called "ARCore."

There are now thousands of applications that use ARKit, ranging from apps that let you measure things to apps that radically change the real world in terms of everything from shopping to new forms of video games. If you visit Apple's store at its new headquarters in Cupertino, you will be handed an iPad where you can see an augmented display of the building across the street from the store.

Touch the screen and the roof pops off so you can see what it's like inside. Touch it again and you'll learn about the solar panels that provide clean energy to the building. Touch it again and you can see a wind analysis layer of the building.



Figure 5.14: Photo credit: IKEA. IKEA Place lets you move furniture around in your house and try it out to see if it fits your lifestyle before you buy it and unpack it, only to discover it doesn't work.

It is apps like these that show the power of AR, but they also show you the limitations, of which there are many. Let's dig into a few of the limitations of mobile AR:

- **You have to hold the phone:** This keeps you from being able to use more capable controllers, or even your hands, to touch, manipulate, or shoot virtual items.
- **Tracking often doesn't work well:** This is especially the case in dark rooms or places that don't have a lot of detail or in moving vehicles like planes or buses, and tracking is almost impossible to get. Apple, even in its headquarters, put a grid pattern on top of the building model so that the camera on the iPad would be able to figure out where it was and align virtual items accurately and quickly.

- **Occlusion doesn't always work:** Occlusion means that virtual items properly cover up things behind them. Or, if a human or, say, a cat walks in front of the virtual item, the virtual item disappears behind that moving human or animal properly.
- **Virtual items don't look natural for a variety of other reasons:** It's hard for your iPhone to be able to properly place shadows under virtual items that match real-world lighting conditions. It's also hard to make virtual items detailed enough to seem natural. For things like virtual cartoons, that might not be a big deal, but if you want to "fool" your user into thinking a virtual item is part of the real world, it's hard to do because of the technical limitations inherent in today's phones and the software that puts these virtual items on top of the real world.



Figure 5.15: Photo credit: Niantic. Pokémon GO lets you capture characters on top of the real world in its popular mobile game, which generated about a billion dollars in revenue in 2019.

The other problem that most people don't think about when they are holding their mobile phone is that their screen is actually fairly small. Even the biggest iPhones or Samsung devices are only a few inches across. That hardly is enough to give your brain anything close to the immersion, embodiment, and other advantages of VR. The problem of having to hold the phone or tablet is the most daunting, actually, particularly for workers. Imagine working on a factory line where you need both of your hands to insert a part into a product. Now, imagine having to put down that part to hold up a phone to augment your work and see details on what you need to do next to properly hook that part up.

Or, if you are playing video games, imagine trying to play paintball, or go bowling with your friends. The act of holding a phone, especially with a tiny screen, limits your ability to recreate the real-world version of those games.

This is why companies like Apple, Qualcomm, Magic Leap, Facebook, Google, and others have already spent billions of dollars to develop some form of wearable AR glasses, blades, or head-mounted displays, if you wish to call them that.

Wearable AR

The industry has been attempting to have a go at this for quite some time and already there have been some notable failures.

ODG, Meta (not associated with Facebook), DAQRI, Intel Vault, and others populate the AR “burial zone” and insiders expect more to come. This is a risky game for big companies. Well, to really take on the mass market, a company will need the following:

- **Distribution stores:** The ability to get people to try on these devices.
- **A brand:** Because these computers will be worn on your face, consumers will be far more picky about brands. Brands that are perceived as nerdy will face headwinds. Already, Facebook announced it is working with Luxottica, which owns brands like Ray-Ban, Oakley, and many others.
- **A supply chain:** If their product does prove successful, tens of millions will buy it. There are only a few companies that can make those quantities of products and use the latest miniaturization techniques, along with the latest materials science, to make devices lightweight, strong, and flexible to absorb the blows that will come with regular use.
- **Marketing:** People don’t yet know why they need head-mounted displays, so they will need to be shown the advantages. This will take many ads and other techniques to get consumers to understand why these are better than a smartphone alone.
- **Content:** If you have a pair of \$2,000 glasses and you can’t watch TV or Netflix and play a lot of different games at home, or use them at work with a variety of different systems, you will think you got ripped off. Only a few companies can get developers to build these, along with making the content deals happen.
- **Design and customer experience:** Consumers have, so far, rejected all the attempts that have come before. Why? Because the glasses that are on the market as of early 2025 generally have a ton of design and user experience problems, from being too heavy to not having sharp enough visuals, to looking like an alien force made them.

Only a few companies have the ability to get where consumers want them to be. One company we have heard about has built thousands of different versions of its glasses and it hasn't even shipped yet. The cost of doing that kind of design work runs into hundreds of millions of dollars.

- **Ecosystem:** When we got the new Apple Watch and the new Apple AirPods Pros in late 2019, we noted that turning the knob on the Watch (Apple calls it the Digital Crown) causes the audio on the headphones to be adjusted. This kind of ecosystem integration will prove very difficult for companies that don't have phones, watches, headphones, computers, and TV devices on the market already.
- **Data privacy:** These devices will have access to more data about you than your smartphones do. Way more. We believe consumers are getting astute about the collection of this kind of data, which may include analysis of your face, your vascular system (blood vessels), your eyes and what you are looking at, your voice, your gait and movement, and more. Only a few companies can deal satisfactorily with regulators concerning this data and can afford the security systems to protect this data from getting into the wrong hands.

This is why when Meron Gribetz, the founder of Meta (not associated with Facebook), told us that he thought he had a chance to be disruptive to personal computers, we argued with him. He thought the \$200 million or so that had been invested in his firm would be enough. We saw that a competitive firm, Magic Leap, had already raised more than a billion dollars and thought that meant doom for his firm. In April 2020, Magic Leap announced massive layoffs of about 50 percent of its workforce, or more than 1,000 people, and refocused away from attempting to sell to consumers.

We turned out to be right about Gribetz's Meta, for his firm was later decimated and sold for parts. Others, too, had the same dream of coming out of the woodwork and leaving a stamp on the world. ODG's founder, Ralph Osterhout, had quite a bit of early success in the field. He invented night-vision scope devices for the military and that got him the ability to have a production line and team in downtown San Francisco to come up with new things.

He sold some of his original patents to Microsoft where they were used on the HoloLens product. But, like Gribetz's company, his firm couldn't stay in the market.



Figure 5.16: Photo credit: Robert Scoble. Meron Gribetz, founder of Meta, shows off how his headset could augment the world and let you use your hands to work.

These early attempts, though, were important. They woke developers up to the opportunities that are here if all the pieces can be put together and they saw the power of a new way of working, and a new way of living, thanks to AR.



Figure 5.17: Photo credit: Intel. Intel Vaunt promised great-looking smart glasses with innovative digital displays but didn't provide the features that potential customers would expect.

They weren't alone in their dream of changing computing for all humans, either. Some other notable failures tell us a bit about where we are going in the future. Intel's Vaunt aimed at fixing a ton of problems with the other early attempts. It had no camera due to the PR problems that Google Glass got into and the last version could only show 2D overlay visuals in black and white. It also had no speaker or microphone. It looked like an everyday pair of glasses.

Too Much and Too Little

It was Intel's fear of the social problems of smart devices on your face, though, that doomed it. In April 2018, it announced it was canceling the Vaunt project. To do AR and the virtualized screens that would give these devices real utility beyond just being a bigger display on your face, the device needs to know where it is in the real world.

Not including cameras meant no AR. No microphones meant you couldn't use your voice to control the devices. No speakers meant you still needed to add on a pair of headphones, which increased complexity, and consumers just decided to stay away. Not to mention that Intel didn't own a brand anyone wants on their face and didn't have the marketing muscle or the stores to get people to try them on and get the proper fit and prescription lenses.

DAQRI, on the other hand, went to the other side of the spectrum. They included pretty great AR and had a camera and much better displays. Workers wearing them could use them for hours to, say, inspect oil refinery pipes and equipment. But, they still went out of business.



Figure 5.18: Photo credit: Robert Scoble. DAQRI's headset gathered hundreds of millions in investment and showed promise for enterprise workers, but never took off because of its tether and expensive costs.

Where Intel threw out any features that made the glasses seem weird or antisocial, DAQRI kept them in. It was big, bulky, and had a tether to a processing pack you would keep on your belt or in a backpack. This appeared very capable, and because they were aimed at enterprise users who didn't need to worry about looking cool or having a luxury brand associated with the product, it seemed like it would see more success. For a while, it did, except for one thing: the price tag.

Many corporate employees heard about this system, gave DAQRI a call, and then found out the costs involved, which could run into hundreds of thousands of dollars, maybe even millions, and that killed many projects. In September 2019, it joined the failure pile and shut down most operations.

Getting Past the Setbacks

Now, if we thought the failures would rule the day, we wouldn't have written this book. Here, perspective rules. There were lots of personal computing companies that went out of business before Apple figured out the path to profitability (Apple itself almost went bankrupt after it announced the Macintosh). Luckily, the opportunity to improve how we all compute is still keeping many innovators energized. For instance, at CES 2020, Panasonic grabbed attention with a new pair of glasses. Inside were lenses from Kopin, started by John Fam.



Figure 5.19: Photo credit: Robert Scoble. John Fam of Kopin shows off how his lenses work in a demo. His lenses and displays are now in Panasonic's latest head-mounted displays and soon, he promises, others.

One thing that's held back the industry has been the quality of lenses and displays, and Fam is one of the pioneers trying to fix the problems. What problems? Well, if you wear a Magic Leap 2 or a HoloLens 2, you'll see virtual images that are pretty dark, don't have good contrast, are fairly blurry, and are fairly thick. They also have disadvantages, including a small viewing area, usually around 50 degrees, when compared to the 100 degrees on offer in most VR headsets. Because these new displays and lenses need to display virtual images while letting a user see through to the real world, they also are very difficult to manufacture, which is why a Magic Leap 2 is priced at \$3,300 and a HoloLens 2 is \$3,500. Today, there are only a few approaches to getting virtual images into your eyes. One approach is waveguides. Here, tiny little structures in something that looks like glass bounce light dozens of times from projectors until the right pixel gets to the right place and is reflected into your eyes. You can't really see these extremely small structures with the naked eye, but this process makes for a flat optic that can be fitted into a head-mounted display. It has lots of disadvantages, the most significant of which is that waveguides are typically not very efficient.

The ones in the first HoloLens presented virtual images that could barely be seen outdoors, for instance. They also present eye-strain problems because it is hard to make these kinds of optics present images at various focal depths. Getting virtual images close to your eyes? Forget it, but even ones that are a few feet away don't usually work at various depths as well as taking into account things in the real world, which puts strain on your eyes. Two terms optics experts talk a lot about are vergence and accommodation.

Vergence and Accommodation

Vergence is how your eyes move inward as you try to focus on things that are closer and closer to you. Cross-eyed? That's vergence, and if your glasses don't align images well as they get closer, they can cause your eyes to get much more tired than when they have to focus on real-world items for a long time.

Accommodation refers to how your eye changes shape—and how the lenses in your eyes move to focus on things close up. The eye often uses cues from the real world to figure out how to do that and in smart glasses, these cues aren't available because everything is presented on a single focal plane. Eye strain appears again as an issue.

Microsoft with HoloLens 2, its latest AR device, went in a different direction—it uses a series of lasers, pointed into tiny mirrors that reflect that light into the back of your eye, which allows you to see virtual images. So far, this approach has brought other problems. Some users say they see a flicker, while others say that colors don't look correct.

Either way, we haven't seen one monitor/lens system yet that fixes all these problems. The closest we've seen is one from Lumus, but because it stayed with waveguides, it needed huge projectors pushing light into the optics. This might be good for something as big as a HoloLens (which weighs about a pound) but the industry knows it needs to ship products that are closer to what Intel's Vaunt was trying to do if it wants to be popular with users.

We hear new optics approaches are coming within the next few years, which is what will drive a new wave of devices, but even those spending billions are no longer hyping things up. Meta's founder, Mark Zuckerberg, has been much more muted when discussing the future with the press lately. Zuckerberg now says he believes we will get glasses that replace mobile phones sometime in the next decade.

Fast, Cheap, or Good

The saying goes that when designing a product, you can have it cheap, fast, or good, except you can only have two of the three. By "good," we mean well designed and appealing for the user: a good and intuitive user interface, appealing aesthetics, and so on. If you've got a "good," fast product, it's not going to be cheap. If it's both fast and cheap, it's unlikely that it's had so much time and expertise put into the design process to bring you a "good" user experience. Spatial Computing product designers will have to make other choices on top of this canonical choice as well. Lightweight or the best optics? A cheap price tag or the best performance? Private or advertising-supported? Split the devices up into their component parts and you see the trade-offs that will have to be made to go after various customer contexts. Aiming a product at someone riding a mountain bike will lead to different trade-offs from aiming one at a construction worker inspecting pipes, for instance.

Since we assume you, our reader, will be putting Spatial Computing to work in your business, we recognize you will need to evaluate the different approaches and products that come to market. These are the things you will need to evaluate for your projects. For instance, having the best visuals might be most important to you, and you are willing to go with a bigger, and heavier, device to get them.



Figure 5.20: Photo credit: Robert Scoble. DigiLens display at Photonics West in 2019. Photonics West is where the industry's optics experts go to show off their latest work and here a potential customer tries on DigiLens' latest offering.

Here are the 10 technology areas where devices will compete:

- **The monitors and optics:** These control what you see. The ones that present more data in a bigger field of view with more pixels, or higher resolution, are usually in the bigger and more expensive devices than the ones on the lightweight side of the scale.
- **The GPU:** GPU speed means how many polygons per second can be processed. The more powerful the GPU in your device, the more detailed the virtual items presented can be. Again, the more powerful ones will be in devices that can supply the battery power these will need and the heat handling that comes along with more powerful chips.
- **Battery:** If you want to use your devices for more hours without needing to charge, you'll need a bigger and heavier one. There will, too, be various approaches to charging said batteries that you will need to consider.

- **Wireless capabilities:** Will it have a 5G chip so it can be used on the new high-speed wireless networks? Or will it only have Wi-Fi capabilities, like current Microsoft HoloLens devices? That's great if you'll only use it at home or work, where WiFi abounds, but makes it harder to use elsewhere.
- **The sensors:** Newer and bigger devices will have several cameras and a variety of other sensors—one to look at your mouth for data to build avatars and study your sentiment, one for your eyes, and another to gather light data, so it can adjust the monitor to be brighter or darker depending on conditions, and four for the world, two of which will be black and white for tracking and two will be color for photography and scanning items. In addition, there's usually an **Inertial Measurement Unit (IMU)**, which watches how the unit moves. An IMU actually has a group of sensors—an accelerometer, a gyroscope, a compass, and a magnetometer. These give the system quite a bit of data about how the headset is moving. The top-of-the-line units also have some 3D sensors. All these sensors add many new capabilities, including the ability to have virtualized monitors that stick to surfaces, such as real-world tables or AR characters, masks, items, and such.
- **AI chip:** Newer devices will add on chips to decipher what the unit is seeing and to start predicting the next moves of both the user and others in their field of view.
- **Corrective lenses:** Many users will need custom lenses inserted, or systems that can change the prescription of the system. Apple and others have patented various methods to bend lenses to supply different corrective prescriptions.
- **Audio:** On some systems, we are seeing array microphones in some high-end units (an array has a number of microphones that can be joined together in software to make it so that they work in higher noise situations, such as a factory floor). Also, a variety of speakers are available, with more on the way. Some even vibrate the bones in your face, which passes audio waves to your inner ear.
- **Processing:** Some systems are all self-contained, like the Microsoft HoloLens, where everything is in the unit on your head. Others, like Xreal and Magic Leap, need to be tethered to processor packs you wear, or to your phone.
- **The frame and case:** Some design their products to be resistant to a large amount of abuse, while others design them to be hyper-light and beautiful. There will be a variety of different approaches to water resistance and shock resistance, and trade-offs between expensive new materials that can provide flexibility and reduce device weight while looking amazing, versus cheaper materials that might be easier to replace if broken.

Each product that we'll see on the market will have various capabilities in each of these areas and trade-offs to be made.

Differences to Be Aware of

We won't go into depth on each, since we're not writing this book to help the product designers who work inside Apple, Facebook, or others, but want to point out some differences that business strategists, planners, and product managers will need to be aware of.

For instance, let's look at the lenses. Generally, you'll have to know where you will attempt to use these devices. Will they be outside in bright sunlight? That's tough for most of the current products to do, so might limit your choices. Lumus, an optics company from Israel, showed us some of its lenses that would work outdoors. They said they are capable of about 10,000 nits (a unit of measure for how bright a monitor is), versus the 300 to 500 nits in many current AR headsets, which only work in fairly low light.

The problem is that Lumus' solution was expensive and fairly big. This is appropriate for a device the size of a HoloLens, but you won't see those in glasses anytime soon. The Lumus lenses were a lot sharper than others we had experienced before and since and had a bigger field of view, but product designers haven't yet included those in products, in part because of their big size and higher cost.

At the other end of the scale were lenses from Kopin and DigiLens. DigiLens showed us glasses that were very affordable (devices would be about \$500 retail) but had a tiny field of view—about 15 degrees instead of the 50 that Lumus showed us, which is about what the HoloLens 2 has, and not as much resolution either. So, reading a newspaper or doing advanced AR won't happen with the DigiLens products.

So, let's talk about how these products are made. When we toured Microsoft Research back in 2004, we met Gary Starkweather. He was the inventor of the laser printer. Starkweather talked about his research into the future of computing. Starkweather was looking for ways to build new kinds of monitors that would give us ubiquitous computing. In hindsight, this was the beginning of where the HoloLens came from. In a separate lab at Microsoft Research nearby, Andy Wilson was working on software to "see" hands and gestures we could make to show the computer what to do, among other research he and his team were doing. Now, this technology is embedded in the HoloLens, Meta Quest, and Apple Vision Pro headsets.

He still is working on advancing the state-of-the-art human-machine interfaces at Microsoft Research today. Starkweather, on the other hand, died in 2019, but his work was so important to the field that we wanted to recognize it here.

From Kinect to the HoloLens

This early research from the mid-2000s soon became useful as 3D sensors started coming on the scene. In 2010, Microsoft announced the Kinect for Xbox 360, which used LIDAR technology. This included a 3D sensor that attempted to bring new capabilities to its video game system. The system could now “see” users standing in front of it. Microsoft had licensed that technology from a small, then-unknown Israeli start-up, PrimeSense. This LIDAR technology is now ubiquitous—found on more advanced mobile phones (such as iPhone Pro) and XR headsets.

Why was Kinect important? It soon became clear that Kinect wouldn’t be a success in the video game world, but Microsoft’s researchers saw another way to use that 3D sensor technology. Shrink it down, put it on your face, and aim it at the room, instead of players. That would let a wearable computer map the room out, adding new computing capabilities to the room itself.

That’s exactly what Alex Kipman and his team did on the first HoloLens. That device had four cameras and one of those 3D depth sensors that mapped out the room and then converted that map into a sheet of polygons (little triangles). You can still see that sheet of triangles once in a while when using the HoloLens and needing to do calibration.



Figure 5.21: Photo credit: Microsoft. HoloLens 2 brings a much more comfortable headset with way better gesture recognition, along with a flip-up screen so workers don't need to remove it as often.

The sensors on the front of a Spatial Computing device, like a Microsoft HoloLens, build a virtual layer on top of the real world where virtual items can be placed. Software developers soon become adept at manipulating this virtual layer that many call the “Ar Cloud.” This database of trillions of virtual triangles laid on top of the real world provides the foundation for everything to come, from simple text to fully augmented or virtual characters running around.

When the HoloLens was introduced, it brought a new computing paradigm. It’s useful to look at just some of the new things it did that weren’t possible before:

- **Screens can now be virtualized:** Before HoloLens, if you wanted a new computer screen, you had to buy one from a retailer and hook it up. They are big, awkward, heavy, and expensive. With HoloLens, you just gestured with your hands or told the system to give you a new screen. Then, you could grab that virtual screen and place it anywhere in the room around you. There aren’t any limits on how many screens you could get, either.
- **Surfaces can be changed:** When using RoboRaid, one of the games we bought for our HoloLens, aliens could blow holes in our real walls and then crawl through them to attack. The effect was quite stunning and showed developers that you could “mix” reality in a whole new way, keeping the real world but replacing parts of it with virtual pieces.
- **Everything is 3D (and not like those glasses at the movies):** The first thing you figure out when you get a HoloLens is that you can drop holograms, which are 3D items that move, nearly everywhere. We have a clown riding a bicycle around our feet as we write this paragraph, for instance. Now, that might not sound that interesting, but you can view your business data in three dimensions too. Over at BadVR, a start-up in Santa Monica, California, Suzie Borders, the founder, showed us quite a few things, from shopping malls where you could see traffic data on the actual floor to factories where you could see data streaming from every machine. This new way of visualizing the world maps to the human brain and how it sees the real world, so it has deep implications for people who do complex jobs, such as surgeons, for instance.
- **Everything is remembered:** Our little bicycle-riding hologram will still be here even a year later, still riding in the same spot; the monitors we placed around our office stick in the same place. The data we leave around stays where we left it. Microsoft calls these “Azure Spatial Anchors” and has given developers new capabilities regarding this new idea.
- **Voice first:** The HoloLens has four microphones, arranged in a way to make it easy for it to hear your commands. You can talk to the HoloLens and ask it to do a variety of things and it’ll do them and talk back. While you can do the same with Siri on an iPhone, this was built into the system at a deep level and feels more natural because it’s on your head and always there.

- **Gesture- and hand-based systems:** Ever wanted to punch your computer? Well, HoloLens almost understands that. Instead, you open your hand and a variety of new capabilities pop up out of thin air. You can also grab things with your hands, and poke at them now too, all thanks to the sensors on the front of the device that are watching your hands, and the AI chip that makes sense of it all.

Given this combination of technologies and capabilities, which are superior to those that the Magic Leap or any other current AR headset has, we believe that the HoloLens historically will be viewed as a hugely important product that shows the world of Spatial Computing at a depth we hadn't seen before.

Technology Takes a Leap

All others would use the HoloLens as a measuring stick, which brings us to Magic Leap. Magic Leap gathered investors with a promise that it would make a big magic leap into this new kind of computing. In hindsight, we don't think the investors knew how advanced Microsoft's efforts were. Its first round of \$540 million came back in 2014, two years before HoloLens hit the market.



Figure 5.22: Photo credit: Magic Leap. This is the Magic Leap 2

The hype before Magic Leap shipped its first product, the ML1, was that it constituted a real graphics breakthrough, with optics that were way ahead of others in the market. That didn't come true, at least not in its first product.

However, the second version, the Magic Leap 2, which is geared toward the enterprise, is a much better headset. There was a lot of turmoil at Magic Leap prior to the release of Magic Leap 2 and there was speculation that the company would shut down. It was able to get more funding and has ended up surviving.



Figure 5.23: Photo credit: XREAL. The XREAL Air 2 is priced at \$399.

This brings us to XREAL and a lawsuit that Magic Leap filed against that company.

Patents and Problems

XREAL (formerly Nreal) started when Chi Xu left Magic Leap in 2016, disappointed with the speed at which the well-funded company was moving. He, along with Bing Xiao, co-founder and chief optical engineer, and Chi Xu, the founder, joined forces in China. The subtext is that China has long been seen as a place that plays loose and fast with intellectual property designed in the United States. Magic Leap promptly sued XREAL and accused Chi Xu of stealing AR secrets from his former employer. It was alleged that Magic Leap was using patents it acquired from Meta in its complaint against XREAL, and Rony Abovitz tells us that intellectual property rights are important to Magic Leap, which has invested billions in its product development. This case will play itself out, but for now, XREAL is getting products on the market that are fairly capable and at a lot lower cost than Magic Leap—it was selling a \$499 model like hotcakes at CES 2020 and now they have entry-level AR glasses, XREAL Air, priced at \$299. Having the Chinese supply chain at their back door will prove to be a big competitive advantage for XREAL, we think, and make it a company to watch.

This leads us to the elephant in the room: Apple. Apple has all the ingredients to make the meal already, from distribution to the brand to the supply chain with Apple Vision Pro.

Compare Apple with XREAL and others and you see the problem. While glasses from XREAL have pretty good optics with bright images and a wide field of view, the company doesn't yet do very much that will get consumers excited, and it isn't clear how it will excite developers enough until it sells millions of units. In Silicon Valley, this is seen as a "chicken and egg" problem—how will it sell enough without apps, and how will it get apps without selling enough?

That all said, we have seen inside R&D labs and know that there are many products under development and that by 2028, we should see massive numbers of people wearing Spatial Computing devices. That will open up many new business opportunities and that's got us excited. For instance, there will be a hunger for new kinds of content; so, we'll see new kinds of cameras and other technologies to help satiate that hunger.

Apple Vision Pro

Ever since we have had our hands on Apple Vision Pro since it became available in February 2024, we have been impressed. That's not to say that there are no negatives. But for a first-generation product, it's pretty stunning.

Want to watch a movie as if you were in a movie theater with superb audio? Do it with Apple Vision Pro. Want to work with as many screens as you want to and use your eyes, hands, and voice to navigate? Want to have a video call with 3D personas and share your work while also tweaking it on the side at the same time? Do it with Apple Vision Pro.

What are the negatives? The headset is a bit too heavy and there have been reports of people developing headaches. The battery power can run anywhere between two-to-three hours before you have to take off the headset to recharge (and that takes a while). The price is very high, starting at \$3,499. We understand it is a first-generation product, but really?

We are dedicating a whole chapter (*Chapter 13*) to the impact of Apple Vision Pro on Spatial Computing so we will go into much more detail on Apple Vision Pro there.



Figure 5.24: Photo credit: Apple. Apple Vision Pro.

Volumetric and Light Field – Capturing Our World in a New Way

DARPA had a challenge: can you take a photo through a bush?

That led to computer vision breakthroughs at Stanford University and other places that are still changing how we look at capturing images today. That challenge, back in 2007, caused Stanford University researcher, Marc Levoy, to move away from using a single lens to take a photo and, instead, build a grid of cameras, all connected to a computer where the software would gather rays of light coming through a bush and collect them piece by piece from all the different cameras, creating a light field. Then, a computer would re-assemble those pieces, sort of like putting together a puzzle.

Shrinking Concepts, Expanding Horizons

Today, the light field concept has shrunk from using individual cameras arranged on a set of shelves to using pieces of image sensors, sometimes even putting microscopic lenses on top, and the field of computer vision has greatly expanded. Computer vision is used in a variety of products, from a June Oven, which can tell whether you have put a piece of toast or a steak into your oven, to autonomous cars, where computer vision is used to do a variety of tasks.

The technique gathers data from a variety of sensors and fuses that into a “frame,” which in AI lingo is all the data from one specific point in time. A Tesla’s frame, for instance, includes the data from its seven cameras, its half-dozen ultrasonic sensors, its radar sensor, and a variety of other IMUs that have motion, compass, and other data. Once a frame has been captured, then a variety of computer vision techniques, which are a subset of AI, can be performed on that frame.

These new techniques have deep implications for entertainment, and, well, pretty much everything—since we’ll soon see camera arrays in tons of products.

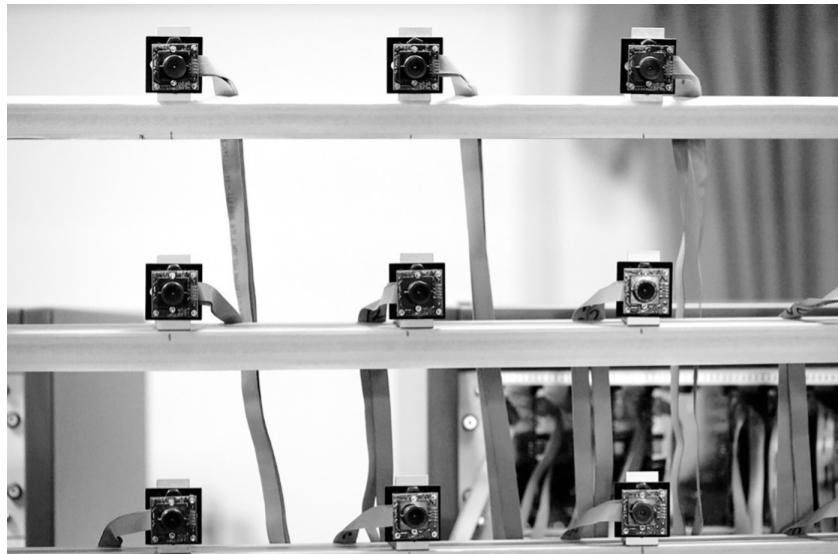


Figure 5.25: Photo credit: Robert Scoble. Part of Marc Levoy’s photographic research lab at Stanford University in 2007.

This research led to the development of new kinds of cameras that are still being improved today.

Already, camera arrays are being used to capture dancers on “So You Think You Can Dance” and many other things, particularly important in Spatial Computing’s VR and AR. When Buzz Aldrin talks to you in a VR headset, he is captured on a camera array with a new technique called “volumetric” video, which means you can walk around Buzz’s virtual image as he talks to you.

Shooting Some Volumetric

Before we go on, we should define some terms. Volumetric video does what it sounds like: it measures volumes. When creative people say, “I’m shooting some volumetric,” that means they are capturing video inside a special studio that has one of these arrays of cameras.

We visited one multimillion-dollar studio, which had been owned by Intel, near Hollywood, to get a look at how it's done. There, a sizable room had a dome surrounding a space that could fit half a basketball court. In the dome, as we looked up, we could see more than 100 cameras.

The walls of the dome were green so that computers could get rid of the dome and easily insert virtual imagery later in editing. Under the dome, Intel had more than 100 cameras connected to a data center where the videos from each camera were stitched together to make something that editors could use and then distribute to people who would see that scene in a VR headset. This 3D dataset is actually pretty useful for regular movies and TV shows, too, because there are infinite camera positions surrounding the action. So, a director could choose a different camera angle to include in the movie they are directing. If that one doesn't work, they could choose another of the same scene. There's a whole industry of volumetric studios, including Metastage, which was backed by Microsoft and utilizes Microsoft technology for volumetric capture for VR and AR experiences.



Figure 5.26: Photo credit: Intel Studios. This volumetric capture dome had 100 4K cameras around the dome, all pointed inside. When joined together, the cameras make a 3D image set that you can walk around, or through.

This is different from 360-degree video. Where volumetric studios include hundreds of cameras aimed inward, toward a single point, 360-degree cameras start at a single point and aim outward. This is a useful technique to capture, say, video of yourself skiing down a hill, since when you get home you want to be able to see the mountain in all directions, including your family members skiing in front of and behind you as you glide down the mountain.

This brings us to light fields. Strict volumetric cameras don't capture the data about all the photons of light streaming off of the subjects that are being captured. Light fields go further, to actually record the angles that light is reflecting off of subjects and entering a camera array.

This, in theory, will let you capture scenes in a much more natural way, which looks amazing to your eye if you view it in a Spatial Computing device on your face.

The real magic of light fields, though, will be that they will enable far more interactive scenes and much greater editing control over objects and scenes. Today, a 360-degree video, and even most volumetric data, is not interactive at all. You sit there and you watch, or with volumetric video, you walk around while watching. But soon, in scenes with light field capture and lenses, you might be able to walk around things and even touch them and pick them up (given enough programming to enable that). Light fields are the holy grail of capturing light, as they do so in such a detailed way that you can do a ton more with that data later.

Video professionals are pushing even further into this field of using arrays of cameras to do new things. Michael Mansouri, founder of Radiant Images, a firm that used to rent cameras to movie studios and now does R&D on cameras for those studios, showed us his latest camera arrays that are doing things Ansel Adams never dreamed of doing with a camera.

Camera Magic

Mansouri's cameras can capture thousands of frames per second, for instance, and are synced with extremely fast strobe lights—this super-fast capture lets him build massive datasets that contain all the data to build light fields or volumetric constructions. He arranges about 100 of these cameras around an actor to capture him or her in ways that just weren't possible a few years ago. His cameras capture every ray of light and the angle at which it entered the camera array, so computers later can "remix" the data.

He said that he is doing tons of tricks to capture more data than the human eye can see, too.

For instance, TV images run at 30 frames per second. He can push his cameras to capture many more frames than is needed to satisfy the human eye and he oversamples those frames—the cameras he is using in one of his arrays can capture thousands of frames per second, which makes for great slow-motion captures, but he is using fast-frame captures to do new things and capture new data that helps AI do new things with video. One frame might be overexposed, another underexposed. One might be black and white, another with color saturation pushed high. Then, the computer vision systems that he and others are building can find new details in the data to make images not only sharper and clearer but also add capabilities that are hard to explain in a book.

The effect of all this light field work will be a video that you could walk around in a very natural way. Imagine walking through a meadow in Yosemite at sunset. It won't be like a stage, where the facade is exposed as soon as you walk behind something and see that it's just a board painted to look nice from the front; here, you can walk around every plant and tree, and the deer 20 feet away looks real. You'll never find yourself "backstage."

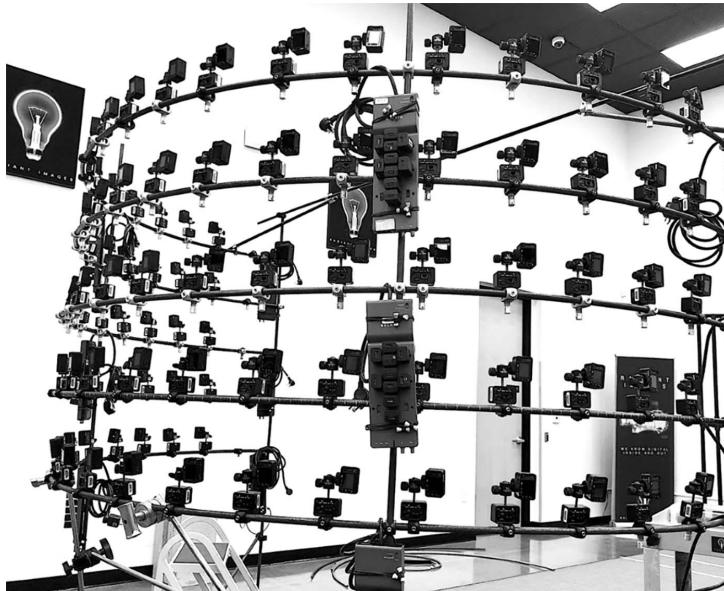


Figure 5.27: Photo credit: Radiant Images. Radiant Images' Nebula Light Field and volumetric system has dozens of high-frame-rate video cameras that it uses to both capture volumetric and light field data.

That kind of experience will arrive soon to our headsets. Michael says, "Why stop there?" He sees a world where you will walk onto a virtual football field and watch the Super Bowl from a new vantage point: that of the quarterback actually playing the game. His team actually shot a commercial for Oakley that gave a taste of how that would feel.

What comes after that might blow even him away: new datasets from the real world might let you walk off of that virtual football field and back into your seat, or even to places outside of the stadium, with lots of virtual holograms along the way. A decade from now, Hollywood could use these new datasets to bring us an entirely new form of entertainment. One that isn't delivered on film anymore, but, rather, a dataset we call an AR cloud that streams to your headset as you walk around.

AR Clouds (3D Mapping) and Internet Infrastructure Including 5G

As we move into a world of glasses, we will want to move away from a world that requires so many apps. After all, do you want to load an app just to see where a product is when you get to a shopping mall? Or, do you want to load another if you get to a ski resort just to see which runs are open? Or, yet another to see the menu at a restaurant? No.

What strategists are planning is a new contextual system that brings you functionality where and when you need it, and that stretches onto the field of entertainment as we just discussed.

A New Kind of Map

When you walk into a shopping mall, you'll see all sorts of virtual helpers. When you go skiing, you'll see the names of the trails as you ski past. When you go to a grocery store, assistants will pop up, showing you more details on top of products, along with navigation to the products that are on your shopping list. When you drive, you'll see other things, like maps and controls for your car. When you exercise, you'll see your health apps contextually pop up. This contextual system will also let you experience new forms of entertainment, from football to movies, where you can follow an actor into, say, a burning building.

Now, some of this can happen with simple location data, but in a true Spatial Computing world, there will be virtual beings moving around, helping you out, and you will want to play new kinds of games against other people, along with tons of augmented, or automated, things that either add data onto the world or keep track of the world to add new utility.

To get to this deeper augmentation, literally every centimeter in the world will need to be mapped in a new 3D dataset that most people haven't yet considered. These maps barely exist today, mostly being built for autonomous vehicles or robots to use to navigate around streets. The Mapbox chair of the board and former CEO, Eric Gundersen, laid out how this will all work. His firm already provides maps to millions of developers and if you use apps like Yelp, Snapchat, or Foursquare, you are already seeing maps provided by his firm.

He and others see a need for a new kind of map: one that doesn't just have a single line for a street, but that looks like the real world.

Mapping cars are building the skeleton for this new kind of map, which some of the industry has started calling an "AR cloud." As vehicles roll by on the street, they are making high-resolution maps and ingesting, via AI, every sign and lots of other features.

If you see Apple or Google's latest maps, you'll notice they are getting more 3D features every day, and, in the case of Google, are able to use AR and computer vision to sense where you are, thanks to you turning on your camera, and show you navigation aids. This is using an early form of this kind of AR cloud and gives us a taste of the world that will rapidly evolve over the next decade. Google is actually using their captured data from their mapping technology Street View to deliver a product called Google Geospatial Creator, to be used along with ARCore and Adobe Aero.

Some places, though, will need to be mapped out in a lot more detail. For instance, if you want to play games where monsters jump around your home, like you can with a HoloLens 2 or Magic Leap 2, then those systems will need to build a 3D map of your walls, ceiling, floor, and any objects in your home.

Building a Map

When the Niantic-acquired 6D.ai used the term AR cloud (it was the firm that popularized the term), that's what they meant. That firm had up until recently shipped a **Software Development Kit (SDK)** that had let developers build AR cloud capabilities into their apps.

What did 6D.ai's code do? It asked users to turn on their camera and then, as users moved their phone around where they were, it instantly built a 3D structure of what it saw. Moving around enough, it soon captured a digital twin of your living room. We've used this software in parks, stores, our yards, and at shopping malls. In each space, it built a new 3D map.

Now, this map, which usually is kept hidden from users, is quite special. It is actually a sheet of polygons, those little triangles that are built by computers to simplify a 3D world to a dataset that a small computer can keep track of.

As Matt Miesnieks, the then co-founder and CEO of 6D.ai, showed us the app, we noticed that it was building a virtual copy of the real world—all made from those polygons—and then all sorts of things could be done with that virtual copy. He showed us virtual things moving around that virtual room and new kinds of games where balls bounce off of these virtual walls. Keep in mind that this “copy” of the real world is laid on top of the real world, so the user thinks the balls are bouncing off of their real walls. They aren't; the code for the game is built to work with the AR cloud, or digital copy, of the house that was just scanned.

Soon, technology like that built by 6D.ai will map out almost everything as we walk around, or, at least, use maps that either a mapping car or someone previously walking around built for us. Often, these new 3D infrastructures won't be seen by users. After all, that's the case already.

Do you really know how Google Maps works under the little blue dot that shows your phone on a street? Very few do.

We will share a bit about how Spatial Computing will work (autonomous cars, robots, and virtual beings will all need these new maps) because there aren't any standards today. Mapbox's data doesn't interoperate with the maps that Apple or Google are building, and we can list dozens of companies that are building 3D maps for various reasons, from Tesla to Amazon, along with a raft of others that aren't well known by consumers, like Here Technologies or OpenStreetMap.

If you think about having a world where hundreds of millions, or even billions, of people are walking around with glasses and other devices that are using these maps, you start to realize something else: current LTE wireless technology, which delivers 2 to 10 Mbps of bandwidth, just won't be able to keep up with the data demands of having to download and upload trillions of these polygons every time someone is moving around the world, not to mention all the other data that users will embed into this AR cloud. We're going to talk about Gaussian splats next, a very new technology that actually does not use polygons to map out the world.

3D Gaussian Splats

3D Gaussian splats are a method used to display and work with data in three dimensions in a smooth, graduated way. This technique is commonly employed in fields such as computer graphics and data visualization for a variety of purposes.

This is really early, but we see a future where 3D Gaussian splats could replace volumetric imaging. Right now, if you take several photos from different viewpoints, the technology joins those viewpoints together to get a 3D Gaussian splat image.

Here's a simplified explanation.

Understanding 3D Data

Imagine data that occupies a 3D space. Each piece of this data has a specific location identified by three values: X, Y, and Z. These points could represent anything from positions in a computer-generated scene to data points in a scientific research visualization.

What Does Gaussian Mean?

The term "Gaussian" is derived from the name of the mathematician Carl Friedrich Gauss. In the realm of 3D Gaussian splats, it particularly refers to the Gaussian distribution, famously known as the bell curve. This function shows a symmetrical distribution of values around a mean (average), creating a smooth, bell-shaped curve.

When applied around each data point, it generates a gradient that resembles a softly sloping hill, highest at the center and gradually tapering off.

The Concept of Splat

In visualization, “splatting” refers to the action of spreading out data points over a surface or throughout a space, akin to flinging a small droplet of paint onto a surface, where it flattens and spreads out. For 3D Gaussian splats, each point in the dataset acts as one of these droplets. However, instead of spreading randomly, it spreads in a controlled manner dictated by the Gaussian distribution, forming a smooth, hill-like spread around itself.

Combining the Concepts

Putting it all together, 3D Gaussian splats involve each point in a 3D dataset generating a smooth, hill-shaped gradient around itself, according to the bell curve. This creates a continuous, smooth representation of 3D data, useful in visualizing phenomena like clouds or smoke in a visually appealing manner, or in scientific visualizations to emphasize concentrations or intensities of data points without abrupt boundaries or distinct separations. We see a time when 3D Gaussian splats could be used by Hollywood. But not yet – if you start to zoom in on a 3D Gaussian splatted image, you can actually see the “splats.” Also, you can’t “walk” into a splat, just around it (the same as a volumetric image).

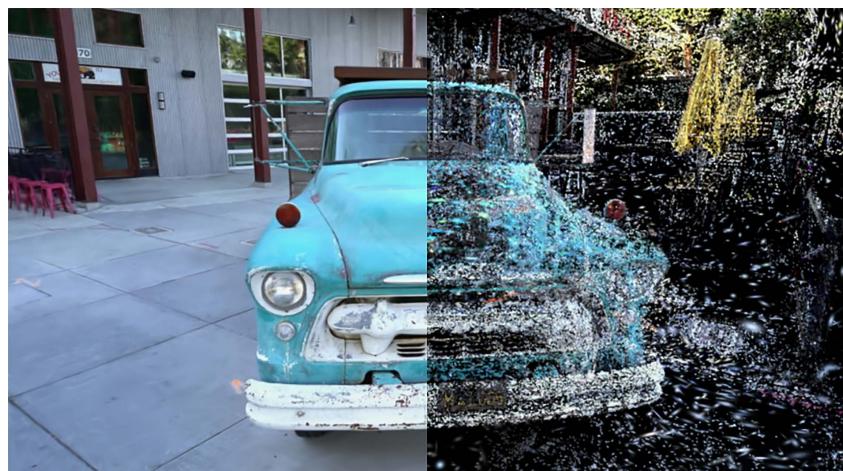


Figure 5.28: Photo credit: Doru Arfire. A 3D Gaussian splat image on the left-hand side. On the right-hand side, a much closer view of what Gaussian splats really look like.

Ready for 5G

Now, you understand why so many of us, like telecom analyst Anshel Sag, are excited by 5G. He explains that 5G is a spectrum of new wireless technologies that, at the low end, will increase our bandwidth more than three times and at the high end, provided by millimeter-wave radio, will give us hundreds of times, maybe even thousands of times, more bandwidth.

Today's 4K videos don't need that kind of bandwidth, but when you have a team of people playing a new kind of virtual football on a field made up of high-resolution AR cloud data, you can see where 5G will be needed.

5G promises not just a lot more bandwidth, which can support these new use cases, but much lower latency too: it should take around two milliseconds to get a packet from your device to a nearby cell tower. That's way less than current technology. That means your virtual football will be much more in "real time" and the ball won't hang in the air while it is waiting for a computer somewhere else to render it on the way into your hands from your friend, who might be playing from a few miles away, or even somewhere else in the world, thousands of miles away.

There's a third benefit to 5G, too, where many times more devices can use a single cell tower. When we go to baseball games with 40,000 other people, it's typical that we can't even make a call because the wireless systems used just weren't designed for thousands of devices to be used at the same time on one cell tower. 5G fixes that problem, and also means that cities and factory owners, not to mention companies like Fitbit or Tesla, with their millions of users and products in the real world, can have millions of sensors, all communicating with new clouds.

The problem is that there are some downsides to 5G. To get the highest speeds (some journalists are showing that they can get around two gigabits per second, which is faster than a fiber line to your home), you need to be fairly close to a cell tower that has the newer millimeter-wave antennas on it (very few have these radios right now) and this new high-spectrum radio wave doesn't go through walls or other obstacles very well; so, even if you live or work near one of these new towers, you will probably need to buy a repeater or mesh network for your home or business to bring that bandwidth indoors. These downsides only apply to the high end of 5G, though. Most smartphone users won't be able to tell the difference anyway. Already on mobile phones, you can stream video just fine with only LTE, and any form of 5G will be just fine on a smartphone. It is when we move to new Spatial Computing glasses that you notice the difference when you are connected to one of the high-end 5G networks, which provides more than a gigabit in bandwidth.

Now, when you see the new iPhone, you'll know that it is really laying the groundwork for true Spatial Computing, which will provide us with capabilities that are inaccessible on our mobile phones, such as interactive virtual football.

New Audio Capabilities Arrive

We are seeing an audio revolution unfold on two fronts: new “voice-first” technology that powers systems such as Siri or Amazon’s Alexa, and new spatial audio capabilities with greatly improved microphones, processing, speakers, and the integration of these things into contextual systems. We are getting tastes of this revolution in Apple’s very popular AirPods Pro.

Audio Augmentation

Apple uses the processing capabilities in AirPods (there’s more compute power inside these little \$250 devices than was shipped on the iPhone 4) to do a new audio trick: pass through real-world sound in a “transparent” mode. That gives you a hint at the potential augmentation that could be done with new audio devices. Apple and others have patents that go further; using bone transmission from devices touching your face, they promise private communication that only a wearer can hear, along with even better audio quality.



Figure 5.29: Photo credit: Apple. Apple AirPods Pro earbuds not only do noise canceling but can mix audio from the real world and digital sources together.

These revolutions in voice and audio are hitting at almost the same time and once we get them built into our Spatial Computing glasses, we’ll see a new set of capabilities that business people will want to exploit.

These new capabilities have been brewing for some time—it's just now that computing has shrunk and microphones have gotten inexpensive. Thanks to the smartphone and its billions in sales, we are seeing these new capabilities coming to consumers.

It was back in 2005 that we first saw array microphones on desks inside Microsoft. Back then, putting four high-quality microphones, and the computer to join those together to do new kinds of sensing, into a box cost about \$10,000. Today, a \$150 pair of Pioneer Rayz headphones have six microphones embedded into their earbuds and along the cord that drapes around your neck. Why use these array microphones? The Pioneers have the best noise canceling we have seen, that's why, and having better, low-noise microphone arrays helps the computer audio response systems that are evolving, like Apple's Siri, work a lot better. After all, if Siri can't hear you, it can't answer you.

Spatial Audio

The processing included inside the headphones can do much more, though. Early headphones had no idea where users were looking. Are they walking north, west, east, or south? The headphones had no idea.

Think about if a system needed to get your attention—hearing “free coffee over here” means a lot more if that sound comes from your left or right, no?

This leads us to predict a new audio-based operating system. You might have different audio assistants located in different places around you. Your email assistant's voice might come from two o'clock. Your social network one? Ten o'clock.

These new audio capabilities are also arriving as new AI and much smarter voice assistants are coming. Apple has rebuilt Siri from scratch, we hear. It isn't alone. Amazon has 10,000 people working on Alexa, trying to make it better than similar systems from Google, Facebook, Samsung, and Apple.

In just a few years, it seems like everything has got Amazon Alexa built in. There are microwave ovens, locks, lights, phones, and much more that now have microphones and can do stuff when you say, “Hey, Alexa.”

To those who have experienced Alexa in their homes, this is probably the first taste of ubiquitous computing—computing that's always on, and always waiting for commands. How did this happen?

We see that these “voice-first” capabilities align deeply with Spatial Computing. Already, our Tesla car is getting voice upgrades that let us ask it to take us places. Soon, consumers will expect the same from all of our computing experiences.

One reason we are excited about Spatial Computing is that it will provide a visual layer to these voice-first systems. A major problem with, say, Amazon Alexa, is that you aren't sure what it actually can do. For instance, did you know that the Internet Archive has the entire Grateful Dead recorded history for you to listen to just by talking to Alexa or its competitor, Google Assistant? It's there; the Internet Archive even wrote an Alexa Skill for it, but figuring that out is nearly impossible unless you already know it's there.

Add in a visual component, though, and voice systems get a lot more powerful. Yes, you could look down at your phone for that visual component, but what if you are skiing and want to talk to Siri? "Hey Siri, what run should I take to get to the lodge?" isn't nearly as nice as if Siri could actually show you which way to turn on your glasses' screen, and asking it, "Hey Siri, what can you do with Foursquare?" would lead to a much more satisfying experience if you had some visual displays to see instead of trying to listen to all the commands it could respond to.

Put these all together and soon we'll get a quite different consumer: one that can find out anything nearly instantly just by asking for help and one that will be fed information via their ears in ways that seem impossible today. Where does that matter? Well, gaming is already using spatial audio in VR to make new experiences possible.

Gaming and Social VR Experiences

Ryan Scoble had a girlfriend he met in VR. They played in Rec Room, the "virtual social club," every afternoon. It had bowling, paintball, and a ton of custom games. He had never met her and didn't know where she really lived. He didn't know her real name (her screen name was LOLBIT) but they played together as if she were a 10-year-old living next door.

He isn't alone in this.

A World Without Rules

VR has brought its magic to gaming and even if you aren't a gamer, it's useful to try a variety of games in VR because of what they can teach you about where the user experience is going. These games are increasingly social and collaborative because VR lets you feel the presence of another person, the same way you would if that person was in front of you in real life.



Figure 5.30: Photo credit: Robert Scoble. Mark Zuckerberg, Facebook founder and CEO, talks about video game success stories at the Oculus F8 Connect conference in 2019. Here, he highlighted the success of games like Beat Saber and then announced a new social gaming platform, Horizon.

This is the lesson that Lucas Rizzotto learned. He's a game developer and has developed several VR experiences, including the popular "Where Thoughts Go," which lets you explore a virtual world populated by other people's thoughts. He says there's a generational shift underway between people who play VR games and those who don't. "Kids will talk with their parents," he says, pointing out that lots of weird conversations are soon going to happen due to this shift. "You won't understand, but I was an eggplant," he says a kid might try soon explaining to a real-world family member.

One of the magic things about VR is that you can be embodied in other forms, either animal, human, or vegetable. Anyone who is playing with Snapchat's AR face features, called filters, knows this, but in VR, it goes to the next level. Chris Milk, founder of Within, a company that made augmented and VR experiences and games, showed us how he could make people into a variety of things: birds, apes, or even fish. We watched him demonstrate this in his VR experience, "Life of Us," at Sundance in 2017 where people started flying within half a second of realizing they had wings instead of arms.

The next generation, he says, will grow up in a world where there are no rules. At least not physical ones. Want to jump off that building in VR? Go ahead, there aren't physical consequences like in the real world. Or, if you want to be an eggplant, you can be that too.

The Ripples of the VR Phenomenon

Some other shifts are underway, at least the major part because of gaming on VR:

- **Players expect more than just entertainment:** They want to pay for something they play through and where they come out with new insights about themselves.
- **Content creation will go mainstream:** Digital skills of the future will be artisanal skills. Already, we are seeing a raft of new tools develop, everything from Tilt Brush, which lets people draw whatever they want in 3D space, to modeling tools.
- **People will build their own realities and experiences:** We are seeing this trend well underway in Minecraft and other games.
- **Crafting will be done with digital materials that don't necessarily exist in the real world:** The new tools are letting people create anything their minds can come up with, everything from complex art to entirely new cities. Already, Rec Room has more than a million virtual places to play in, most created by the users themselves. University College London and the Multimedia Anthropology team, led by Raffaella Fryer-Moreira, are building a virtual twin of native Brazilian tribe villages, using Gaussian splatting objects scanned from their original sacred objects.²
- **Gaming is now 360 degrees and much more interactive:** This is much closer to experiential theater, like the influential "Sleep No More," an off-Broadway play in New York where audience members walk around a huge warehouse and actors are all around, performing their lines. While enjoying the play, you can take in the complex theatre set, even picking some things up and discovering what they do.
- **We will be introduced to new, synthetic characters that are far more lifelike than before:** In 2019, Magic Leap already had shown a virtual being, Mica. Mica looked like a human, and she played games with you, among other things, yet she was completely synthetic. Now, with Generative AI, virtual beings are exploding into our culture.
- **Collaboration is much better and deeper in VR:** You can already play basketball and paintball and do a variety of other things together in VR, thanks to Meta's "Horizon," a social VR experience.
- **It's still early:** but there is a possibility of using Gaussian splatting scans to compose VR, and even AR, scenes and experiences.

It is the collaboration part that's most interesting, though, and has certainly caught Mark Zuckerberg's attention. Rec Room has attracted an audience, and Zuckerberg knows that. Horizon seemed squarely aimed at Rec Room, with very similar kinds of games and social experiences, albeit with a promise of better integration into your existing friend pool.

This leads us to a key point: social software is only really fun if you know the people you are playing with. If you are like Ryan Scoble, you can invest time in a game and make some friends there, but most of us just want to get in and do something fun quickly. Up until now, there just weren't enough other people on VR to ensure that you had someone else to play with, and almost certainly none of them would be your real-life friend. That's starting to change now as more people get VR.

Building a VR Community

It has been tough going building a social experience for VR, though. Many studio heads expected VR to sell a lot better than it did. It turns out that this is another chicken-and-egg waiting game: the social aspect really works best if your friends can play with you, but most of your friends don't yet have VR. Lots of early social VR systems sold for very little money, like AltspaceVR, which had announced it was going out of business and then was purchased by Microsoft at what were rumored to be fire-sale prices. Others have recently struggled to keep their doors open, but have had to shut down or sell off, like Sansar, which comes from the company Linden Lab, developers of Second Life. Philip Rosedale was the founder of Linden Lab before leaving that company to start a VR-based company, High Fidelity. Rosedale told us it is hard to get the critical mass to make social experiences work if only a few million headsets have sold.

Even Rec Room acknowledged that the lack of VR users is a problem and it now has apps on iOS and Android so that mobile players can join some of the fun. This is a smart strategy because this helps mobile users realize they are second-class citizens and aren't playing in an immersive, hands-on-the-ball kind of way. This brings us to a major benefit of playing in VR—you are forced to move around, which leads to greater fitness. In some games, like The Climb, you have to move your hands like you are climbing in the real world. We've found ourselves sweating after a few rounds in games like that, and in Rec Room, with its basketball, paintball, bowling, and other movement-activating games, you see this too. Within's "Supernatural," which is an exercise app, complete with volumetrically-captured coaches and a subscription plan that works like a gym, was acquired by Meta in 2023.

Fitness uses of VR have even spread to gyms. A Boise, Idaho-based start-up, Black Box VR, has brought a new kind of weight-training machine to its gyms. Founder Ryan DeLuca showed us how it works: you put on a VR headset and grab two handles at the ends of cables coming out of the machine.

Then, you select which exercise you want, and it walks you through that inside the headset. At the start of each round of exercises, it starts up a game. Keep up your exercise and the game gets more fun and intense, and you gather more points. It was so fun that afterward, we realized we were sore, but at the time we didn't feel like we were working out on the machine.

Cinematic and Interactive Experiences

Outside of gaming and social VR, there are Spatial Computing experiences currently made mostly for the festival circuit.

Festivals That Include VR

The reason why these experiences are shown this way is that the present state of digital distribution of Spatial Computing entertainment experiences is at a very initial unsophisticated state. We won't be focusing on distribution issues here, as it would take us further away from what the focus of this section is.

Cinematic VR experiences are those that are actually 360-degree experiences—that is, they are not actually VR. However, on the festival circuit, the term cinematic VR has taken off; so, we will briefly address it, as well as those non-gaming experiences that are interactive, hence *are* VR. Additionally, there has recently been a very small and slowly growing body of AR experiences that are narrative in nature, shown at festivals.



Figure 5.31: Photo credit: SXSW. People at SXSW 2019 enjoying cinematic VR in Positron Voyager chairs.

Major festivals where cinematic VR, interactive VR, and AR experiences are shown include the Sundance Film Festival in Park City, Utah; SXSW in Austin, Texas; the Tribeca Film Festival in New York City; the Cannes Film Festival (with a specific category for XR experiences added in 2024); the Venice Film Festival; and the Vancouver International Film Festival. The Toronto International Film Festival did show cinematic VR for a few years but discontinued it after having a change in leadership.

There has been much talk about how Spatial Computing festival experiences have been overly experimental and avant-garde and not made with the public in mind. We tend to agree with this characterization of the state of affairs. One of the major issues that impacts what kind of content gets made is how much money is made available to produce it. From 2015 until 2017, many VR headset companies, including Facebook, Samsung, and HTC, provided content funds to jump-start the industry. However, ever since 2018, much of this money has dried up.

Outside of self-financing narrative content for Spatial Computing environments (VCs and other kinds of investors do not generally have content funds), it is very difficult to have the resources necessary to make these kinds of experiences. As a result, only those who are extremely driven and also have access to funds are able to create them. The kinds of narratives that get developed, as a result, are not necessarily those that could actually be marketed for revenue. This brings us back to the original problem: that there is no real place to market Spatial Computing narrative experiences for profit.

Content Production and Previsualization

Over a decade ago, while working on *Avatar*, the filmmaker James Cameron created a technique where actors wear motion-capture suits while being filmed inside digital backgrounds in real time. In films such as *Ready Player One* and *Solo: A Star Wars Story*, among others, filmmakers have started using VR assets and headsets to plan shots using a virtual world set. Cameron also used VR for the previsualization of his *Avatar* sequels.

Another filmmaker, Jon Favreau, has been especially enthusiastic about the use of VR for content production and the previsualization of scenes. His most recent film where he utilized VR in this way is 2019's *The Lion King*, and he actively used his technique, called "V-cam," for episodes of *The Mandalorian*, a Star Wars series for the Disney+ streaming service.

Favreau also used a more nascent version of his technique in 2016's *The Jungle Book*. Additionally, he partnered up with start-up Wevr to make *Gnomes and Goblins*, an interactive VR fantasy world.

The V-cam technology helps with fast turnarounds on schedule and budget, allowing VFX supervisors to create the look of scenes in real time and providing major benefits to the actors and the crew, effectively, according to Favreau, turning the film's production process into a "multiplayer filmmaking game."

Except for a single photographed shot, *The Lion King* was filmed entirely in VR. In order to work out scenes, Favreau and others would put on their VR headsets to plan exactly where their cameras and lights should go, using handheld controllers to move the virtual equipment to create the scenes using the decided-upon calibrations.

Regarding the use of VR for previsualization, Favreau has also said, "We'll probably have to come up with some sort of new language... It's nice to be able to turn to these new technologies that could otherwise be a threat and use them to reinvent and innovate."

The future of the creation of movies and episodes points to actors in VR headsets doing their scenes inside a virtual setting, captured by virtual cameras controlled by a VR headset-wearing crew.

The Merging of the Virtual and the Real

What we've been showing you is how a new metaverse will appear. An entire set of virtual worlds, or layers, that you can experience with Spatial Computing glasses will soon be here. As you look around with these Spatial Computing glasses, you will see the real world and these virtual layers "mixed" with each other in very imaginative ways, and we just presented the pieces that are currently under development that will make this happen. Often, the changes will be small, maybe a virtual sign on top of food, but other times, they might be huge, like presenting a complete virtual world you can move around in and interact with. Over the next few years, we'll see an acceleration toward these new virtualized worlds as the devices that we wear get far more capable and lighter, and new experiences are created for them.

In our next chapter, we discuss the benefits that Spatial Computing brings to manufacturing. The use cases there have a much more practical nature than what we discussed in this chapter, though Spatial Computing's abilities will similarly change how we work.

References

1. You can see the 'Project Glass: Live Demo at Google I/O' video here: <https://www.youtube.com/watch?v=D7TB8b2t3QE>
2. The virtual twin of native Brazilian tribe villages: <https://www.uclmal.com/virtual-museum>

6

Vision Three: Augmented Manufacturing

We had the privilege of visiting Tesla's factory a couple of times. Over the space of just one year, we observed some radical changes. There are more robots, sensors, and screens, as well as closer working spaces with hundreds, if not thousands, of workers sitting around the factory looking at computers running things.

If you are paying attention online, you will see the hype about the changes: "Industrial Revolution 5.0," the headlines scream. The last two decades have seen much digitizing on the factory floor. That change isn't done yet; we still see lots of non-digital processes in many factories that have yet to change from paper and pen to computers.

That said, some major new trends have shown up. We haven't yet seen a world where a lot of workers are wearing augmented or Virtual Reality devices on their faces. Those will come over the next few years, and we talked with manufacturing innovators about where things are going and what changes they are planning.

The changes that are coming are extreme. A new kind of manufacturing is coming: an almost completely automated factory floor, as is already running in China and elsewhere. We'll see more and more AI every year with the addition of new humanoid robots and new AI-driven approaches.

Robots and Other Augmentation

The Tesla factory, for instance, has 1,200 robots, and that's up from prior years. Tesla isn't alone. Many companies are seeing more robots in their manufacturing lines. The word from Elon Musk, Tesla's CEO, is that Tesla's humanoid robot, Optimus¹, had started working on the factory floor by the end of 2024.

It used to be that robots were things kept in cages. These robots were too dangerous for humans to get close to. They moved very heavy parts quickly. When we visited Ford's Rouge Factory, which builds a new F-150 truck every minute, we saw such a robot installing windshields into new trucks.

While many robots continue to work away from humans, a new kind of robot is showing up—"cobots." These new robots work right next to humans, and even with humans aiding them in tasks. Controlling these new cobots is difficult, though. After all, it's easy to tell another human "Hey, can you help me hold this plate while I screw it down?" but it's pretty hard to do that with a robot, especially if you have to touch a tablet or a mouse hooked up to a PC sitting on the workstation you are at. Soon, those who work in factories tell us, you'll be using AR glasses with microphones and 3D sensors and software that can see what you are working on, who with, where your hands are, and where your eyes are looking, and then the "cobot" revolution will emerge in a big way on factory floors.

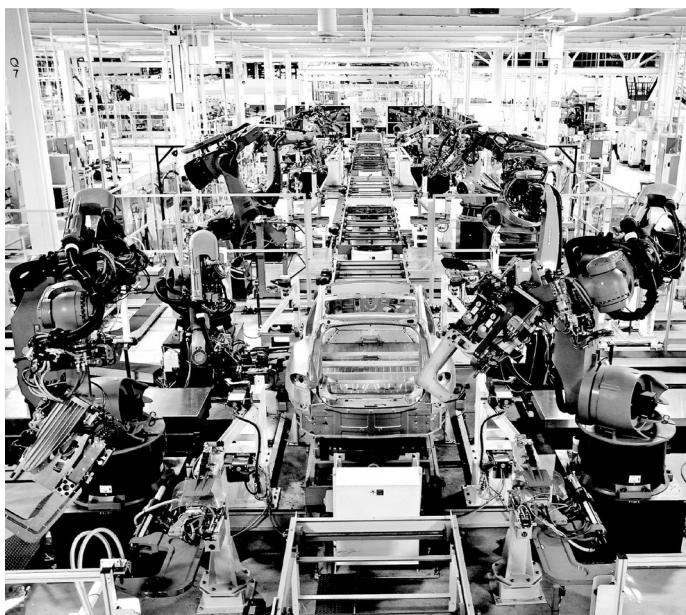


Figure 6.1: Photo credit: Steve Jurvetson. Tesla factory filled with robots building the Model 3.

As we talked with factory workers across different factories, we heard of a clash between using old and new technologies and systems. We define “old” as not using cameras or 3D image sensors to watch you work, worrying about the impact of privacy on workers, and not wearing AR headsets or glasses to assist you in working. Also, “old” is not having robots integrated into your manufacturing lines working with humans, not using additive manufacturing or 3D printing to make parts, and not being able to radically change a factory floor within days (Tesla set up an entire line in a tent, for instance, within a few weeks to make a version of its Model 3 car).

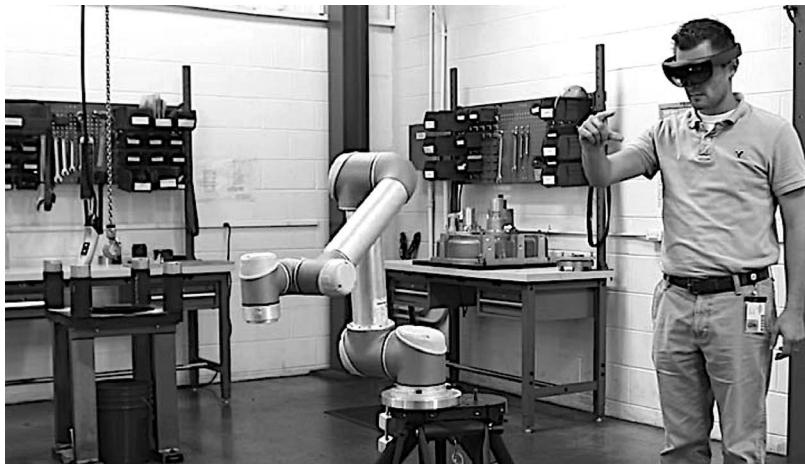


Figure 6.2: Photo credit: Universal Robots. HoloLens and cobot interaction testing.

Obviously, talking about the status quo in such a way opens a very controversial set of subjects: from the protection of jobs to privacy at work. Different companies and different cultures will come at these things with different sets of constraints. We have already found differences in how factory workers look at cameras on factory floors, for instance. Those in Germany are far more conservative than those in China. The same goes for workstations or glasses that “watch” worker performance. At Jabil in Silicon Valley, a company that makes many manufacturing workstations around the world, we saw AR workstations that told people how to build items but also watched and graded workers at the same time.

Dennis Morrow, former UX Design Manager at Radius Innovation, a subsidiary of Jabil, told us they were moving those augmented displays into headsets so that the worker is augmented all the time, not just while sitting at a workstation. Such a set of displays on a worker’s face could even direct them to the shortest line in the cafeteria. When we visited factories in China, which often have tens of thousands—if not hundreds of thousands—of workers, simply moving people around more efficiently can create sizeable productivity gains.

When talking to innovators in factories around the world, we keep hearing about “Industrial Revolution 5.0.” Industrial Revolution 5.0 is all about how AI fundamentally changes how things are done, taking into account how humans are using it as a tool, including the human and environmental perspective². Industrial Revolution 4.0 was all about digital twins, and now, with 5.0, it is about how humans actually purposely work to make manufacturing happen—how humans and machines work together.

AI, robotics, IoT, autonomous vehicles, 3D printing, nanotechnology, biotechnology, materials science, energy storage, quantum computing, and augmented and Virtual Reality are all exponentially evolving and are bringing the perfect storm of changes to literally every part of human life. But how we design and make things may see the biggest changes, and, certainly, have deep implications for employment and what jobs look like.

A future car company might ship a potential customer a pair of Augmented Reality glasses before the car is even introduced. Those glasses may let you see the experience of owning one, let you order one, and let you watch it being built in real time. Think that’s impossible? Well, we are seeing a world by 2030 where Augmented Reality glasses will be fairly low cost, subsidized by advertising. We are reminded that Tesla doesn’t spend much on advertising and we talked with hundreds of people standing in line overnight to buy a Model 3 without having seen the car yet.

Tesla, in a way, provides great insight into how quickly work is changing. We visited its factory in Fremont, California, several times over the years, and each time we visited, we noticed more robots, more sensors, more computers, and a higher density of workers working with robots.



Figure 6.3: Photo credit: BMW. Using a HoloLens to put together a BMW car engine.

In talking to innovators who are designing new factories like the ones Tesla runs, we found that they are planning to introduce new technologies at a rapid rate.

Volkswagen, for instance, has a complete virtual copy of its factory floor. It isn't alone; we heard this from many manufacturers we interviewed, from Coca-Cola to Audi. This "digital twin" lets workers redesign factory workstations, get trained on new jobs, and soon, even control machines and processes on the factory floor from anywhere—not to mention get remote assistance from vendors and other employees. The digital twin can be walked around using virtual or Augmented Reality headsets and is being made smarter every day.

We heard of many who are letting workers make digital documents, videos, audio recordings, and 3D objects accessible via their digital twins so that future workers would get the benefit of past learning. At many factories, managers are worried about a new "silver tsunami." That is, older workers who are retiring in droves, leaving the factory floor with dozens of years of knowledge about how things are run. New workers are currently learning from those workers digitally while navigating to their workstations. The digital twin can be overlaid on top of the real factory floor and can be used as the basis for adding all sorts of new capabilities both to factories and to workers wearing new Spatial Computing devices.



Figure 6.4: Photo credit: Volkswagen. Digital twinning at the Virtual Engineering Lab of Volkswagen Group IT.

Just one of these new capabilities, remote assistance, doesn't sound that significant, but when you consider Boeing's factory in Everett, Washington, where they build 747, 777, and 787 airplanes, it takes nearly 30 minutes to walk from one end of the factory to the other.

The building is so large it has its own weather pattern inside. If technology can save a worker a trip, you're saving the company money and making the worker happier.

That same factory has workers wearing devices that "augment" their work, showing them where one of the thousands of wires in an average airplane should be plugged into thanks to a camera on the front of the device that sees a barcode on each wire and on each piece of equipment that a wire needs to be connected to. Think through the trends hitting consumer expectations. Things will be done with higher quality, fewer materials, lower cost, and more choices—all faster.

All of those expectations are forcing factory owners to constantly add new technologies to make things cheaper, better, faster, and with more transparency. Political concerns and human rights are driving many of these. Want to make sure children aren't making your shoes? We can see future shoes coming with a video of every step of their manufacture to prove that. Why will that be possible, if not probable, in the 2020s? When we toured Tesla's factory, we met some LG employees who make cameras, and they said that cameras are getting cheaper and with a higher resolution because of the billions of dollars being spent in R&D by the smartphone industry. We are seeing this dramatic decrease in the price of technology hitting nearly every product, from the original June Ovens that have cameras inside to Monopoly board games that now have voice commands thanks to a little computer that ships with its latest game.



Figure 6.5: Photo Credit: Boeing. Using HoloLens to view a wiring application displaying digital engineering drawings inside a Boeing airplane.

Where robots inside factories used to be big, expensive machines that were kept separate from humans, in cages, robots now are integrated across the factory floor and are delivering parts to workstations in such a way that they navigate around humans walking the corridors. Spatial Computing technologies are making those robots “see” and are giving them much better navigation skills, just like the same technologies being used in autonomous vehicles on the road.

Ensuring really deep integration between human workers and robots, though, requires new AR glasses that haven’t shipped yet. These glasses will allow a worker to, say, install an air conditioning pump in a new car while talking to and controlling a robot that’s holding that pump. Today’s computers that need to be touched, or wheeled over to, just aren’t up to the job.

The iPhone currently has a sensor that can “see” over 300,000 points of data. That gives us a hint as to what’s coming to the manufacturing floor. Those sensors won’t be limited to just iPhones. Microsoft, with its Azure Kinect sensor/camera, can already be used to do new things on factory floors, from watching product quality to building new systems for robots to do new work. We are told by many innovators that computer vision will be used across the factory floor and will provide data to know where everyone and every part is in real time.

These techniques are already being used while building planes at Boeing, but the cost of the systems means that they are used in complex situations, though the cost has been going down over the last few years. We expect that as costs come down further, thanks in great part to including them in mobile phones, we’ll see them used in much more routine work where the return on investment isn’t nearly as high as figuring out which of the 20,000 wires needs to be plugged into a certain outlet inside a Boeing plane.

We’ve laid out the context somewhat. Manufacturing is becoming more hypervigilant. Hypervigilance means that fewer mistakes will be made, and work will become more efficient because the technologies used are more detail-oriented, digital, and data-based. Jabil’s Dennis Morrow told us of a new kind of screwdriver that is used to put in screws so small that the worker (and the robots who will also use it) needs to work under a microscope. We visited a watch factory in Switzerland where \$100,000 watches are hand-built using similarly small screwdrivers and devices that magnify the work. “Today a user has to look on a 2D screen to see that screw,” he said, while pointing out that they are building an AR system that lets users zoom in without looking away from the work. It goes further than just magnification, though. It shows the right process, which screws to put in and in which order, then checks the work while it’s being done, warning you if the work is being done wrong. “How can we empower them to have these kinds of superpowers at their fingertips?”

This new system, driven by data and AI, with visuals in AR devices on workers' faces means a huge improvement in quality, Morrow said. It also means that QA people down the line have fewer things to reject and fix.

Where this becomes quite evident is a company called Lumafield. Lumafield makes large refrigerator-sized CT scanners that can be used by manufacturers to find out if a product they are making has any defects. It could also be used to recognize when a product is a forgery. The CT scans create 3D images that could then be examined in an AR, VR, or MR headset, although the image could be viewed without one. Looking at these images with a headset, however, would provide immensely more benefit and ease.

AI is used to initially discern the differences between non-deficient and deficient items during the scan and also could be used after the scan to further differentiate between the two.

Workers who would normally be in charge of recognizing deficiencies would be replaced by Lumafield's large CT scanners. The only thing that would be left to do is to examine the scans.

We talked with Eduardo Torrealba, CEO of Lumafield, who told us that his scanners cost \$75,000, much less than other manufacturers of large scanners that could go for \$250,000 to \$500,000. Lumafield's scanners are made to scan products that are football-sized or smaller, while the more expensive ones are made for the medical field. Therefore, Lumafield does not have any direct competitors (yet).

Yes, there will be job shifts, many factory managers told us, but they see a new opportunity for workers to be rapidly trained for newer, more interesting jobs due to this hypervigilant manufacturing capability. Some of the major new Spatial Computing tools that will enable workers and factories to become more productive and efficient will now be discussed further. One of those tools is the use of AR.

Augmented Reality Visions and Informed Robots

Within a decade, working in a factory will probably be very different from how it is today. Many, if not most, workers will probably be wearing AR glasses. They will be working with a lot more robots and will be guided by a connected system that will track everything, including machines, humans, robots, and parts flowing into the factory. Additionally, production itself might be quite different thanks to additive manufacturing (3D printing).

Let's go back to the new devices being worn on the face by many workers. We call them Spatial Computing visors, but others call them AR or MR glasses or visors or **Head Mounted Displays (HMDs)**. These new displays let you see computing in 3D and all around you.

They also include microphones, 3D sensors for your hands, for your eyes, and for the world around the worker, along with sensors inside the devices that track up and down, left and right movements, and other things, from a potential heartbeat to the brightness of the room.

On little displays, workers see 3D and 2D images overlaid on top of the real world, or, as optics get better, replacing pieces of the real world. Bruce Dickinson, Vice President and General Manager of Boeing's 767/747 program, commented about Boeing's use of AR in their factories: "The cross-functional team working on this technology has made a step-change breakthrough in our quality and productivity by following their passion to pursue a great idea. We don't often see 40% improvements in productivity, and I'm convinced that it was a culture of innovation and leaders who are willing to say 'yes' that enabled this idea to come to life."

While most of today's Spatial Computing devices are big and have relatively low-resolution monitors (which is why you don't see many workers using them for coding or writing reports, for instance), in a decade, they will be small and have a much larger field of view, with sensors that will see 100x more points of data than they do today, and with GPUs that are much more powerful.

They will also be connected with wireless systems that'll be 10x or more faster than what we can get today. These glasses will also have eye sensors and voice recognition systems that make the Siris and Alexas of 2024 seem pretty stupid indeed, along with hyper-sharp optics that will replace all sorts of physical screens.

"Applying computer vision and AR tools to the manufacturing process can significantly increase the speed and efficiency related to manufacturing and in particular to the manufacturing of automobile parts and vehicles," Tesla's application for a patent for a future AR device reads.

Elon Musk, Tesla's CEO, noted in an interview with Marques Brownlee, technology journalist/YouTube star, that many jobs are still not appropriate for robots because they are too complex or still have too much variability. He pointed to a harness of wires being strung along the roof and sides of the car. "If a robot faults out, you have to have a 24/7 technician," he said, while pointing out that such a robot failure had him woken up at 2 a.m.

Much work is being done by robots, but they need to be programmed and dealt with as machines, and not in a very human way. In a decade, robots will be responsive and much more like humans, thanks to much better voice commands, but the worker will be able to interact with the machines around them in a much better way because the glasses that worker is wearing will let the worker point, grab, and control the robot virtually and probably even remotely.

Giovanni Caturano, former CEO of SpinVector, spins it around (he builds AR and VR software for factories). He thinks the robots will train us humans on what they need through these glasses. He sees that AI can learn new things faster than we can, which is why AI is beating us at Go and other games but lacks lots of other skills, or is unable to physically reach and grab something. So, the AI needs to ask us for help, he says, and that's the kind of system he's already working on.

Think about moving a heavy box with a friend. There's a lot of visual information you both share. Is one side of the box going lower? Is your friend's face straining? Don't you watch for gestures? Listen to voice commands? Yes. Today's factory robots have the beginnings of such human interaction styles, but they will pale in comparison to tomorrow's robots and factory automation types. Caturano explained that robots can watch humans far better than humans can watch robots. Additionally, they won't get distracted, they won't get scared and drop an expensive part, they can see in all directions, and they are better at seeing very small changes in temperature, angles, distances, or pressure than humans are.

To really understand the changes, we talked with Rebecca Johnson and Bettina Maisch of Siemens. Johnson is head of the research group and builds UX prototypes and platforms there and is the CT senior key expert for ambient intelligence.

Maisch was the Portfolio Project Manager for Industrial Design Thinking. Together, they painted a wild picture of the factories of the future.

Maisch noted that such tools will become "exponential learners" in that they will be hooked up to a cloud system that will watch everything about their use, and any learning will be pushed back out to not only the humans but to the networked robots "learning" while using those tools. The whole factory system will become more reliant on data and AI systems that will make that data useful.

Maisch and Johnson surmised that some jobs of the future, both for AI systems as well as humans, will be dealing with what goes wrong. Humans are pretty good about keeping things moving, even when all else goes wrong. One worker who left Tesla wrote about duct-taping a part together to keep the line running. Not exactly good for quality scores, for sure, but if your line is building a new \$60,000 car every 60 seconds, would you stop the line for a part that could be fixed later? Maisch said one of the attributes humans have is working around rules and figuring out how to keep things moving, where robots might just get stuck.

Some of the factory thinkers we talked to even expect that many factory workers will be able to work, at least some of the time, from home, controlling robots from their living rooms—maybe on call just to fix those weird problems that crop up in factories from time to time, whether it's a belt on a machine breaking, or some traffic accident holding up deliveries of parts on the free-ways outside of the factory. But that's a bit too sci-fi even for us. We just see that factory floors are about to radically change thanks to a raft of new technologies and that these technologies will make it a lot easier to retrain robots to do new things and make it a lot easier to work next to, and with, robots.

What we do see are a whole lot of companies building humanoid robots that will be able to work very competently in the factory, as well as the home. These humanoid robots built by Tesla, Boston Dynamics, Figure, 1X, and Mentre Robotics will soon (possibly as early as the end of 2024) be able to do many tasks. They will be equipped with “brains” that use **Large Language Models (LLMs)** to “think” and will be able to understand commands as well as speak. These robots will use computer vision to understand their environment.



Figure 6.6: Photo Credit: Tesla. Tesla's Optimus Gen 2 shows off its dexterity.

One problem with trying to predict exactly what will happen in factories as more robots and Spatial Computing technology show up is that not every factory will get them at the same rate. We learned that in several interviews. Some factories will cling to the old ways of doing things for years, while some, particularly in, say, the automotive industry, will advance much more quickly.

Audi, for instance, already uses AR devices on its factory floors to assist workers with a variety of things. Ashim Guha, former Vice President of Consulting and Enterprise Solutions at Tech Mahindra, told us that AR is already being used quite a bit in Germany and China on such lines, while admitting that other factories he visits are using the same techniques they've used for decades.

On the other side of the spectrum of factories are ones that will be completely automated with no humans in them. Already, there are “lights-out” factories, particularly in warehousing and picking. “Lights out” means there is no need to keep the lights on for human beings.

Fuji Automatic Numerical Control (FANUC), one of the largest makers of industrial robots in the world, has been operating a “lights-out” factory since 2001, where robots build other robots completely unsupervised for nearly a month at a time.

The focus of this section is in the middle of the spectrum: where humans will need to continue to work alongside robots, but where there is enough force to get workers to wear AR devices. That's already happening at Audi, Boeing, and other factories—usually, where there is a complex enough job, like keeping track of tens of thousands of wires inside a plane and where they need to go or be plugged in. A human can do that job better than a robot, at least so far, but the human wears a device like Google Glass that shows a barcode on each wire and displays where that wire needs to go.

In the past few years, though, a new kind of work has arrived thanks to the cobots that we already discussed. These cobots, as we'll see, will get more intimate with humans as we get better spatial tools over our eyes to be able to communicate with them and program them.

It's hard to understand how deeply this new worker will be “augmented” because, well, when we visited factories in 2018 and 2019, whether Boeing, Ford, Tesla, or even Louisville Slugger, we saw very few workers wearing computers on their faces and most robots were still things kept in cages because they were too dangerous for humans to be next to. One set of memorable robots at Tesla that lives in such a cage lifts an entire Model 3 in the air, spins it around, and places it on a new track. This is something that, if a human were in the way, would be quite dangerous indeed.

This new way of 3D working goes beyond design and to the work itself. Maisch told us her team built an app that lets a worker wearing a Microsoft HoloLens work underneath a train engine and it warns that worker if they try to plug something into the wrong socket. Even before that, AR was changing how things are done. If you buy tools for the factory floor or robots, you now can pick them out of a virtual catalog and see how they would work alongside the other robots or tools you have already purchased for your manufacturing line.

The complexity of thousands of robots and humans working together in a modern factory is something that will need new user interfaces and new approaches too. BadVR's founder, Suzanne Borders, showed us a virtual factory floor where you could watch it run from anywhere in the world and see all sorts of data streaming off of the machines.

In the old days, a part might travel around a floor with a simple paper trail kept on a board. Modern factories are moving toward a computing system, one might say an “air traffic control” system, keeping track of literally everything on the floor, and even keeping track of those things outside of the factory, such as incoming component shipments.

Maisch and Johnson said that having new kinds of wearable devices will help workers not just see where the load of parts they need to build something is but also build a new kind of interface with the robots and the factory floor.

Several factory workers we talked with said that the need to bring new tools, particularly AR devices, is coming from two trends: first, younger workers are staying away from factories that don't use them. Second, older workers are retiring and leaving factories with decades of knowledge about how they work. So, many factories are bringing this technology onto the shop floor to record older workers, getting them to leave their knowledge directly on the shop floor through videos, drawings, and 3D captures of what they actually do.

Then the factory managers appeal to younger workers by showing them how useful the new tools are and how exciting working in such a “smart” factory is.



Figure 6.7: Photo credit: Microsoft. ZF Group, a German car-parts manufacturer, is using Microsoft Layout (<https://aka.ms/Esw053>) to help them lay out equipment on its factory floor.

As a part of a “smart” factory, wearables could bring some great benefits. By putting sensors, computers, and displays on workers (or similar tech into robots), and by giving them hands-free interaction, in terms of its benefits, “it’s very convincing,” according to Caturano. The bigger payoff, though, he says, is the lack of disruption and the ability to do quicker upgrades: “The client (a big soft drink company) liked this approach because they didn’t need to shut down the factory for days to wire up lots of sensors.” Also, workers can charge up batteries every night, whereas if sensors were built onto factory floors, new electrical wiring might need to be run, or those sensors would need to run for a year without having batteries changed, he says. New battery configurations are coming in the mid-2020s that will make the issues around batteries much better, and as such, they will revolutionize the way robots, drones, and untethered devices are used for manufacturing and also outside of manufacturing.

Morrow says that the real advantage to all of this, though, is that factories will get rid of 2D screens and move to a system that makes both a human and a robot smarter at the same time.

What really makes a robot “smart” is the combination of computer vision and AI in the form of machine learning software. Computer vision has to do with systems that enable a computer to “understand” digital images, videos, and live-streaming visuals. Part of that understanding has to do with matching and recognizing shapes, colors, and textures.

Machine learning software enables computer vision systems to do that. Before a system can do recognition matching, it has to be trained on many thousands of image datasets. Supervised, unsupervised, and reinforcement learning, as described earlier in this book, as well as other forms of learning, could all play a part in making a robot “smart.” The predominant kinds of features/areas surmised through machine learning include object classification, object identification, object verification, object detection, object landmark detection, object segmentation, and object recognition. Frank Loydl, CIO at Audi AG, emphasized: “Artificial intelligence and machine learning are key technologies for the future at Audi. With their help, we will continue to sustainably drive the digital transformation of the company.”

As expertly explained by Everette Phillips, former Director, North America, of Lyric Robot, one of the major issues, though, when it comes to creating a “smart” robot is that the software that includes both the regular instructional and machine learning parts almost always has to be tailor-made, specifically taking into account the robot and the situational environment the robot is in. Some robots come with ready-made software but then have to be majorly modified, while others do not. In both of these cases, in-house engineers and/or outside consultants have to be utilized to modify and extend existing software or completely build custom versions. An improvement from the past that could be helpful in dealing with this issue is that general manufacturing processes overall have become more flexible so that a robot system might also be partially accommodated by a change in process and even a change in factory floor layout. Beforehand, both the change in process and factory floor layout could be virtually tested by using AR software. The case for the use of Spatial Computing in manufacturing has never been clearer.

According to Deloitte, by 2028, as many as 2.4 million manufacturing jobs may go unfilled, causing \$454 billion in production to be at risk. The main reasons for this are potential workers lacking relevant skills and the retirement of current workers. Spatial Computing could be used to re-skill current workers and then place them in jobs where they work alongside robots.

Arnold Kravitz, former Chief Technology Officer at **Advanced Robotics for Manufacturing (ARM)**, told us that there are novel tests that are being developed for the joint use of cobots and workers using AR glasses that measure and determine how the two could most efficiently “learn” from each other and, once those configurations are discovered, processes could then be updated to take these into account. Similarly, the concept of a virtual digital twin, mentioned previously in this chapter, could be instituted to figure out the best product and manufacturing design and also to find out what the culprits are when a process or system fails.

The Mirror World Comes to Factory Floors

What really is going on is that there's a new boss of the factory. It used to be that there might be someone like Henry Ford walking around, watching everyone working, prodding and pushing the workforce, keeping track of parts that hadn't shown up yet, and so on, but those days are long gone.

The boss of the factory is being replaced by data and AI, along with AR glasses or visors. These cloud computing systems keep track of every little thing in their surroundings. This will likely include even facial recognition to ensure that only the correct workers get access to the factory floor, and the position of the wearable device, down to fractions of a millimeter, along with everything that said wearable device's 3D sensor is shining on.

This new boss is displayed in these Spatial Computing devices as a "digital twin." On one level, the digital twin is just a virtual copy of the factory floor. Every display, lever, switch, motor, and part is there. This makes it really easy to train new workers. You just take them into the digital twin, and they learn how to do their jobs on the virtual copy of the factory floor.

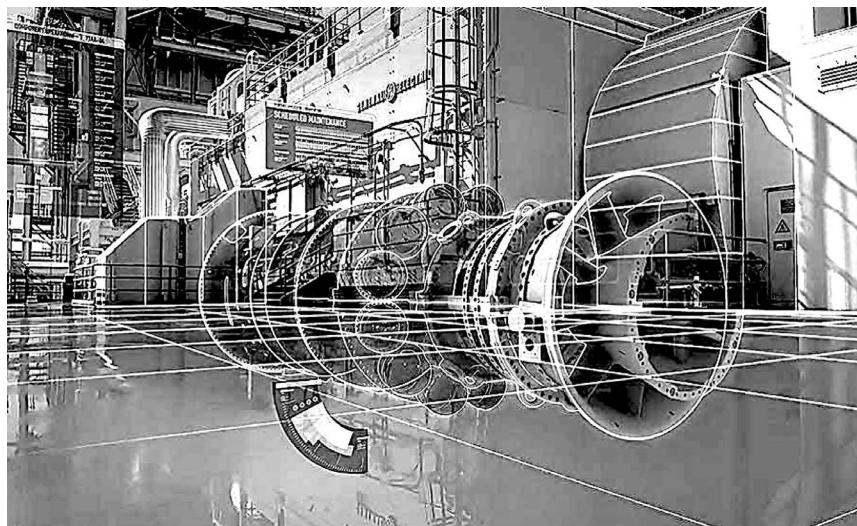


Figure 6.8: Photo credit: General Electric. General Electric is developing digital twins of major portions of its factories and products built to actively aid in machine planning, implementation, and error recognition.

If that's how you look at a digital twin, you'll miss something deeper that's going on: these virtual copies of factories are actually new kinds of databases and virtual interfaces, which physical IoT devices can be hooked into but also other data can be laid onto the virtual floor.

This virtual floor, or rather, digital twin, can then be laid on top of the real factory floor. This new “mirror world” brings radically new capabilities that simply can’t be replicated using older techniques.

For instance, Nucleus VR founder Alexander Bolton showed us how workers can leave training materials, and even instruction manuals and videos, for workers right on top of the real factory machinery. New workers coming into work wearing a Microsoft HoloLens, or another similar device, could “see” the digital instruction manual that got an update since the last shift. They could also see and watch a video left there by the crew that just left, explaining a problem they had with the machinery and how they fixed it.

Also, on top of the digital twin, they could see all sorts of information about the machinery itself. How hot are its engines? How many products a minute are being built? Which machines are running low on supplies? Which ones will need to be taken offline in the next few hours for some maintenance?

Remote workers are even seen moving around the digital twin, and others can be called in with just a command (think about “Hey Siri” or “OK Google” and how that impels a system that is listening for those commands to listen for other commands).

This new digital twin doesn’t just affect the operation of the floor, though. It can save millions of dollars when designing the factory floor. Caturano told us that they now start building the digital twin first. In fact, the digital twin can be designed through AI in a process called “generative design.” Autodesk walked us through this and so did Siemens. A computer will give hundreds, or even thousands, of possible configurations and a human will pick the best one, or refactor and try again. This “generative design” process is already being used by architects. This process would be best done in AR glasses where a factory designer could walk through a virtual factory and see iteration after iteration of factory designs and pick the best ideas that AI comes up with. One example of how this can impact a company is discussed below.

Caturano explained how one problem came up in a factory before they got these new 3D tools to design the factory floor. He told us they thought they had designed it all well but there was a pesky post in the middle of one factory. Yes, the product and the carriages they were being moved around in (think of robot trolleys that can move around, say, soft drinks in pallets) could technically fit, which is why the design got through approvals. But, every once in a while, one of these carriages, when turning, would hit the post, damaging thousands of dollars of product and making a mess. It turned out that moving that post would cost a lot of money. This problem, Caturano says, won’t happen again now that they are designing factory floors on a Microsoft HoloLens where they have a 3D “digital twin” to test everything on before they build the factory floor.

Siemens' Maisch, mentioned earlier, told us that her team designed an app for a mining operation. It turned out that workers usually took one path through the mine, but once in a while, molten steel would flow over the path.

Sensors in the machinery would pick up that the path would be unsafe to navigate, and workers' headsets and mobile apps would navigate around the danger.

This "mirror world" has many new superpowers. For instance, on a real factory line, you might need to build 100,000 parts before you'll see your first defect. Caturano says that on a digital twin, you can force virtual defects down the line, which can be used to train all the AI systems to look for these defects and become far better at spotting them than any human could, even with years of experience. The AI could have dealt with hundreds of forced, simulated defects, while even the most experienced of workers may have only encountered a handful of such problems. The digital twin will soon have extraordinary smarts hooked into each polygon of data that's overlaid on top of the real factory floor.

Everette Phillips, former Director, North America, at Lyric Robot, says, "What you are seeing is an evolution of a change in approach where AI is leaking into industrial systems." He says, "What has been improving is the cost of computational power. That allows more sophisticated code to be run without bringing in more sophisticated programmers."



Figure 6.9: Photo credit: Robert Scoble. Apple's new headquarters AR app shows a "digital twin" on top of a real physical model.

To get a taste of what a digital twin can do, you should visit Apple's new headquarters in Cupertino. There, it has built an AR application that includes a digital twin. On top of a physical model, this digital twin shows you how the headquarters works and looks. You can touch the digital twin to get it to show you various new virtual views.

As you walk around the model with an iPad, you can see how the solar system works, and you can remove the roof and walk around the interior of the building. You can even see a simulation of how wind flows over the building.

This is all while dozens of other people walk around you with their own view of the same thing.

It shows the magic that can come from placing a virtual layer on top of a real-life plant. You can see data flows of humans and equipment in new ways. You could simulate what the plant would look like with a new lighting system installed. You could simulate the impacts on getting workers from parking lots if you built a new shuttle system. That and much more, and, of course, if you have this kind of work done, you automatically get much better training systems and are way ahead in terms of enabling new ways of working.

As factories are increasingly virtualized, they could fundamentally change a lot about what we think of as “work.” Imagine having your morning coffee, signing into your factory floor, and watching your robots in your area doing their work. Imagine saying “Hi” virtually to your coworker Joe, then checking up on all your inventories and data streaming from the machines working... even making an adjustment or two to speed or temperature or checking in a new protocol for one of your robots.

As you walk around the virtual factory floor, you will see updated information and even a new menu for the cafeteria—all virtually. And you can help Joe virtually figure out a problem he's having with one of his robots—all while having morning coffee.

The digital twin will enable all this, plus new virtual interfaces for machines (and other humans) that are pretty hard to comprehend today, but Apple's AR demo at its headquarters gives us a little taste of just how much will change in the next decade thanks to digital twins. Digital twins are among the several kinds of innovations happening in augmenting manufacturing. Another one is the use of robots.

The Human Robot

Imagine being able to safely work alongside robots on the factory floor and being able to communicate with them on a human level. Eventually (maybe sooner rather than later), we have been told, robots on the manufacturing floor and elsewhere will be able to display human characteristics.

Siemens' Maisch and Johnson talked very eloquently about how robots could be visibly humanized for the benefit of humans and how Siemens is running tests along the following lines.

Robots will be able to use “facial” expressions that show whether or not what they are doing is going well or if there is a problem³, as well as using colloquial language. An answer to the question of whether or not everything is going well could be, “Not really...” They could also answer in a depressed tone if there is a problem. This would greatly help in assessing if there is an issue more directly than checking robot diagnostics.

Another kind of humanness that could be integrated into robot systems is the ability to create products that very closely mimic human randomness. Currently, it is very easy to tell the difference between a piece of clothing that has been knit by human hands and one that has been made by factory automation. It is also very easy to tell the difference between a piece of furniture that has been fashioned by a master woodworker and one that has been crafted in a factory. In the future, by using AI algorithms that include LLMs, robots will be able to very successfully mimic human randomness in the making of products, making them appear to have been made by human hands.

This is just the start of making a robot display human-like external and internal characteristics. The question of how far this could go is not the focus of this book, but the forward motion in this direction shows the promise of how far the elements of Spatial Computing could be taken for the benefit of humans.

The arc of where manufacturing is going is clear. Where robots could clearly do a better job than human beings, they will replace them. Where productivity is higher when robots and human beings cooperate, they will be kept as a unit. Computer vision and AI will continue to be used by both humans and robots in greater frequency and for wider uses. AR could be used as the bridge between the human and the robot; however, we do see a day when the robot will be much more human, and vice versa.

This chapter has dealt with how manufacturing is in the process of using augmentation in the form of robots, AR, and AI to improve productivity and efficiency. Our next chapter tackles how the retail industry does the same with regard to logistics and presentation to consumers.

References

1. Tesla's humanoid robot, Optimus: <https://www.youtube.com/watch?v=8Q5Lx2NpYTW>
2. Forbes article on Industry 5.0: <https://www.forbes.com/sites/jeroenkraaijenbrink/2022/05/24/what-is-industry-5-0-and-how-it-will-radically-change-your-business-strategy/>
3. Facial expressions of the Ameca robot: <https://www.youtube.com/watch?v=CSpdEc3sTnE>

7

Vision Four: Robot Consumers

Retail still defines the face of America and indeed the whole world, and you can still see this yourself if you visit small towns or big city malls, whether in Chicago or Dubai. Soon, though, that face will radically change as tons of new technology is brought to all parts of the retail world, from robots to Spatial Computing glasses that consumers will use to shop in new ways. Here, we look at both the past and future for lessons about how deeply our culture and businesses are about to change due to Spatial Computing.

What Drives Consumers?

People buy emotionally, not rationally. That is what we learned while working at consumer electronics stores in the 1980s. What do we mean by that? Let's say you are a grandfather who is tasked to buy some headphones for his granddaughter for, say, a Christmas gift. He goes into a store, say a Best Buy, and is overwhelmed with choices. He can't hear well, so he can't tell the difference between, for example, a Sennheiser or a Sony. The salesperson isn't much better at figuring out what will put a smile on the granddaughter's face when she opens her gift, either.

He remembers that he saw Beats on American Idol and that Apple had bought them, so he walks out of the store with some of those. Why? Because he knows that on Christmas morning, if she doesn't like the Beats, he can defend the purchase by saying, "Hey, if they are good enough for American Idol, they are good enough for me," and that, as Apple's premier brand, he feels good about that, too.

Now, those of you who understand headphones will know that Beats are hardly the best-sounding headphones out there. A rational person would do research on audio profiles and buy coldly based on which headphones sounded best for a particular budget, but that's not how most of us buy, which explains why Beats continues to outsell other brands.

As we talk to retailers, we see that they understand this, and are looking to more deeply figure out how to engage your emotions when you shop, which explains a whole raft of changes underway in retail thanks to Spatial Computing technologies. How do they change retail?

Getting Emotional

Well, these technologies are changing the consumer electronics that people want, which you can see by visiting the **Consumer Electronics Show (CES)** in January in Las Vegas, with its huge booths for devices that use Spatial Computing—from Virtual Reality to drones to autonomous cars to devices in your home that track you and products. These technologies are changing everything about the retail supply chain, store, and e-commerce experiences themselves. For the grandfather, though, buying a set of headphones will mean that not only will the headphones be different but the way he buys them will be, and how they get made and delivered to him will be, too.

Augmented Reality is changing the media he sees. Soon, he will be able to watch American Idol on Spatial Computing glasses and walk around the singers and coaches. Even on current TV shows, we are seeing Spatial Computing technology used in shows such as “So You Think You Can Dance.” Now, a few dozen cameras surround dancers so that producers can spin the view all the way around the dancers. Soon, that effect will seem as quaint as black-and-white TV. He’ll be able to touch products he sees, and ask questions like, “Who made these?”, “Alexa, can you put these on my shopping list?”, “Alexa, what’s the price on these?”, or “Alexa, can you deliver these to my car?”.

Speaking of cars, someday in the next decade, he might even ask his car to go pick them up, all thanks to Spatial Computing technologies¹. That is a bit further out, though, from getting the glasses, which will arrive by the mid-2020s. We cover more of those scenarios in *Chapter 4, Vision One: Transportation Automates*, our chapter on the transportation industry vertical.

Now in 3D

You can already get a taste of how the buying experience is changing by using AR apps on your mobile phone. Before the late 2010s, you would have to go to a store to see how, say, a new couch would fit into your home. Now, you can pull up an app and see a three-dimensional version right on top of your real home to see how it fits. You can even now use many apps to measure your room accurately to see whether that new couch would fit. Lots of apps from stores such as IKEA or Amazon already have these features as we write this book, with many more on the way.

We see a deeper meaning to three-dimensional visualization in retail, though: the act of looking at three-dimensional versions of products and experiences at home hits your emotional triggers far better than looking at grids of two-dimensional products.

Amazon and other stores have seen increases in sales by converting their products to three-dimensional visualizations, and newer stores, such as Obsess (<https://obsessar.com/>), take it a lot further. With Obsess, you see sets of new products in a three-dimensional representation of a real home or office. No more grids of product results.

The result sells way better, particularly groupings of products like, say, art that goes with pillows that go with a vase on a coffee table. You can instantly change how a room looks just by clicking a button and then you can see all the products that go together in that look. This dramatically increases engagement, fun, and the ability to sell groupings of products. There's a reason why so many companies are hungry for consumers getting Spatial Computing glasses, and this new kind of "virtual showroom" is a major part of it.



Figure 7.1: Photo credit: Audi. Audi has deployed over 1,000 VR showrooms in dealerships worldwide.

Audi's virtual showroom decreases product returns due to products not fitting right or not matching what the customer saw online. With AR, for instance, you can try shoes out on your own feet and see how they look. That's quite different, both emotionally and in reality, from trying to imagine a two-dimensional shoe and how it might look on your feet.

At Sephora, they went further. Its Virtual Artist product² came out of its R&D labs in San Francisco in 2019 and has been upgraded over the years. It let you try virtual makeup on your own face with your mobile phone. If you liked what you saw, you could purchase the real thing, and they were careful to make sure the colors matched. This let you play around with (and emotionally bond to, retailers hoped) different looks and makeup styles without worrying what others were thinking about you before you bought. It stopped being available for download in 2022 and we're told that a new app should be coming soon that will replace it.

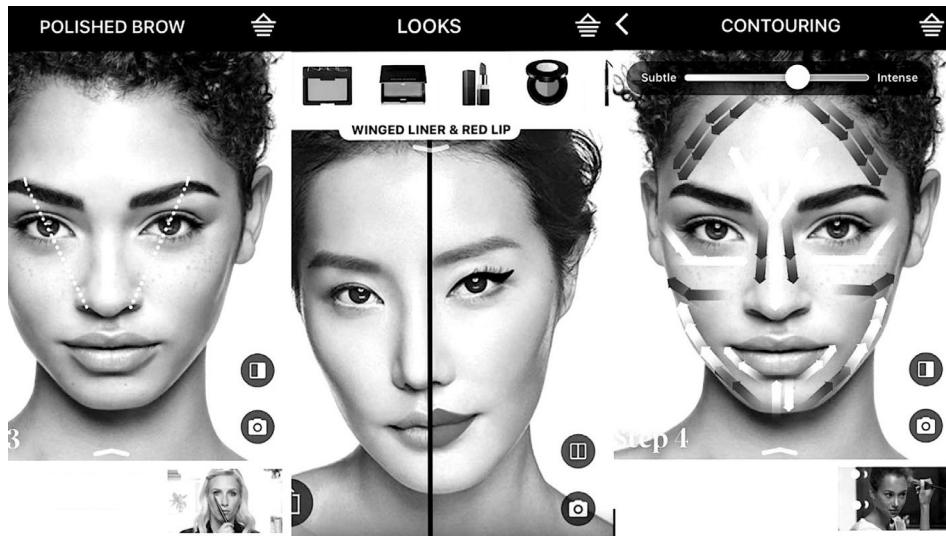


Figure 7.2: Photo credit: Sephora. Illustration of “Sephora Virtual Artist” AR app.

Finally, three-dimensional product visualizations will bring with them a new kind of virtual shopping, one that you can walk through just like you walk through a real-life mall, like the ones in Dubai and Minneapolis we visited to prepare for this book. One thing, though, is that when Spatial Computing devices arrive, you’ll be able to “call up” your friend and walk through the virtual mall with them. Retail is about to see some deep changes due to digital twin technology that we talk about elsewhere in this book, and when we talk to people inside consumer goods companies, such as Red Bull, they are salivating to do new kinds of three-dimensional shopping experiences because they know that these will cause much deeper brand engagement and emotional attachment to brands, if done right.

The visual and emotional side of shopping is just the tip of the iceberg of change that’s coming to retailers in the next decade. Let’s dig into some of the other changes that are already underway with regard to Spatial Computing.

Retail Destruction? No, But Radical Change Underway

During the fall of 2019, one of the biggest Fry’s Electronics locations in Silicon Valley announced it would close. This retail store used to have dozens of checkout lanes and it often had long lines during its heyday in the 1980s. John Fry gave us a tour on the opening day of its Sunnyvale store when it brought in its first million dollars of revenue in one day. Those days are long gone now.



Figure 7.3: Photo credit: Google Maps. The Palo Alto location of Fry's Electronics has closed as of January 2020.

We all know the reason why it closed: Amazon crushed it. The truth is that Fry's didn't change, but the competition did. It was disrupted, as we say in Silicon Valley.

Trying to Keep Up

Amazon and other online retailers, like B&H Photo, took away Fry's unique position of having everything from soldering irons to computer processors. Back then, it was "the place" to buy the stuff you needed to build the products that made Silicon Valley famous in the 1980s as a place to start a technology company. Now, you can get all those online and have them delivered to you, often within an hour or two and for a cheaper price.

Online is not only cheaper, it's better organized, comes with reviews to help you choose between dozens of products, and is better at reminding you of associated things you might also need. One of the last times we visited Fry's in Palo Alto, there was a huge set of shelves with dozens of drones. There was no help to choose. No salespeople nearby. If you did pull one of the drones off of the shelf, you might not realize what batteries you needed to buy, or cases that would work, or even what other devices you would need to go with it. Amazon is different—it shows you all of these things when you buy a drone or, really, any product. The experience of shopping at Fry's is poorer in real life than it is online and online is getting better at an ever-increasing pace thanks to the move to three-dimensional product visualization.

Worse, Fry's didn't have two things that are keeping some brick-and-mortar retailers thriving, like Bass Pro Shops, which sells a variety of outdoor goods, from boats to tents to fishing rods and guns: things you want to try physically before buying, and a showroom attitude that teaches you about the products you are buying. We predict that retailers such as Bass Pro Shops won't disappear in the next decade but can imagine dozens of ways they will be changed as customers adopt more Spatial Computing. We can imagine a virtual stream where you can try out fishing rods and lures and see how they feel, and work, in the stream you are planning on soon fishing in.

This goes way beyond the emotional triggers of changing how we can demo things, now that applications of three-dimensional visualization are becoming more interesting.

We did see the early signs as we visited stores such as Amazon Go (which is now no more), which had no checkout counters. Or in Las Vegas hotels, where small robots deliver things to rooms. Or in warehouses that retailers own, like Kroger, that have hundreds of thousands of robots scurrying around, moving products on their way to both storage shelves and delivery vehicles. Or in the Amazon app, which has a feature where you can get more details by using your camera. Aim your phone at a product and it tells you more. Soon, Computer Vision features like this will bring new capabilities into your lives.

These changes will really be felt in retail stores as consumers move computing from two-dimensional screens to three-dimensional glasses or visors that bring virtualized stores into living rooms and offices, and even if that doesn't happen, three-dimensional product visualization is already changing how people can try on products or see how they fit into their rooms in real life. All of a sudden, you won't need to walk into a store to see displays like you would see in today's shopping malls or grocery stores. These will show up virtually literally anywhere and be much more customizable and personalized.

Looking Ahead

We dream of a world where everyone is wearing Spatial Computing devices on their faces and where Spatial Computing technology also drives around robots, virtual beings, and autonomous vehicles, including drones.

These on-face devices will bring massively new visual experiences into people's homes, and the same devices, when worn by retail workers, can enable new kinds of service, new, more efficient warehouses, and new delivery systems that will get products to you via robots, drones, or autonomous vehicles.

If you have worked in previous versions of retail stores, like we did back in the 1980s while running consumer electronics stores in Silicon Valley, you can see the shifts. Back then, we didn't take inventory on computers. We didn't have cash registers that allowed people to pay via their mobile phones (heck, our customers didn't have computers or phones for the most part either, apart from the rich executives or real estate agents).

Walmart has been training its employees with VR, a trend we expect to greatly accelerate over the next few years because the early examples are proving to not just be profitable but lifesaving (the company claimed that the VR training saved lives during a shooting in one of its stores, which you can read more about at <https://fortune.com/2019/08/20/walmart-ceo-vr-training-helped-save-lives-in-el-paso-shooting/>).

In quite a few warehouses, we are seeing Spatial Computing devices and systems being used to make them more efficient. AR glasses can show you exactly where a product is in a huge warehouse, saving seconds on every trip. Things like dwell time, touching, and transaction data all help manufacturers build products that better serve consumers, and even design packaging that gets their attention on busy store shelves. Soon that will be augmented or, in fact, already is. Walk into a Lego store and you can see three-dimensional models of what is inside each box thanks to an AR system originally built by Metaio, which was later bought by Apple to form the basis of its ARKit platform.

Aspects of retail are changing their roles in several ways due to Spatial Computing. The store itself is morphing into something that serves you in new ways, while augmented stores will let you experience shopping in the comfort of your home in a way that will look and feel pretty close to real stores, and then the retailer can serve you via rolling robots or autonomous vehicles. Let's dig into how the stores themselves are changing due to these new technologies.

Variations on a Retail Theme

The camera functionality in Amazon's app is worthy of checking out and gives a good taste of where things could go in the future. What's funny is that when we ask people if they have tried it, even techie people, very few have. Yet it's been in the app for a couple of years, at least. What does it do? It lets you visually search. Google and others have competitor products, but because it's from Amazon, which has the most highly developed retail store, it works better with the products there.

Aim it at a logo, say the Starbucks logo on your coffee cup, and Amazon will bring up all sorts of merchandise with the same logo on it³. Aim it at a friend wearing, say, a black turtleneck shirt, and it will find you similar shirts to buy. Aim it at a product, or, particularly, a box, like that of a coffee maker, and you will find lots of coffee makers to buy.

When still at Amazon, Chris Hargarten, who ran the largest team of people categorizing all sorts of things, told us they tried to use techniques like these to make the shopping experience even better and faster. “We want users to go directly into the product page they want,” he explained to us. That takes some work, given the millions of products and pages, and the visual search feature inside the Amazon app gives us all another vector to do that.

It isn’t perfect, though. With clothes, it does usually get the style right, but rarely gets the brand right. We hear that over in China, the competitor for Amazon has a similar feature that’s much more accurate. We notice that in China, people use cameras far more frequently than in other countries. This leads us to a new point: if you have more users and even more aggressive users of your technology, you will have more data, and more data leads to more accuracy when you try to do things like figure out the brand of the suit someone is wearing that you like and want to buy for yourself.

The camera, too, could enable a new kind of “showrooming.” This is what they call it when people go into a local retail store, only to go online to buy the product they are checking out for, presumably, a cheaper price. Retailers hate showrooming, and with good reason; Fry’s has been decimated by the practice, and so have lots of other retailers.

But do retailers fight back with technology of their own to help customers build highly customized shopping experiences that lead back to their stores? Very few have the capabilities to hire the software teams that Amazon has. We learned that about 10,000 people are working on the training of Amazon’s Echo line of products alone. What kind of retailer can deal with that kind of technology investment?

Not to mention that Amazon is putting that kind of investment into the entire supply chain. Visit its warehouses and you will find hundreds of thousands of robots scurrying about, moving products from truck to truck for delivery to customers.

Amazon, Hargarten explained, is aiming at making every part of buying and receiving a product more efficient. To see how, let’s walk back into B&H Photo in New York City.

When you walk in there, you have a few choices. Do you just wander around going aisle by aisle, looking to see what is there? Sometimes. We did just that because we enjoy consumer electronics so much, and it reminds us of stores we shopped at, or ran, in the 1980s.

Or, do you know exactly what you want and want to go straight to that section, like we did the time we needed to get another set of Rode Wireless Go microphones? Yeah, there’s even a greeter/security guard at the front door who can direct you right to where you want to go.

The same applies online. Amazon tries to make it very easy for you to go directly to a specific product, and if you don't know exactly what you want yet, it also makes it easy to shop by category, like microphones in the B&H Photo store.

There are points of friction, though. One is that product visualizations are being converted to 3D, but they are stuck on small glass screens where you can't get a satisfying look. Ever try to see how a refrigerator would look in your kitchen by looking at a three-dimensional version on a mobile phone? We have, and the experience isn't satisfying at all. Now that we have the Apple Vision Pro, the situation is changing.

We think the changes are deeper than just trying to bring back brick-and-mortar stores. Everything about retail is changing, from how inventory is ordered to advertising, warehousing, delivery, in-store services, and customer expectations. All of this is being driven by the technologies that drive AR, autonomous cars, or robots. Cameras, Computer Vision, AI/Machine Learning, three-dimensional sensors, and more are changing literally every piece of the retail puzzle. A glimpse of this future can be had if you check out an Amazon Go store.

The Integration of Spatial Computing in Retail—Transforming Customer Experiences, Operations, and Business Strategies

In the ever-evolving landscape of retail, Spatial Computing is emerging as a transformative force, reshaping how businesses engage with customers, optimize operations, and expand their market reach. These immersive technologies—which include game engines and are accessible on the web, on gaming platforms, and through mobile AR or VR devices—have now found their way into mainstream retail, offering a plethora of opportunities for retailers to enhance customer experiences, streamline processes, and explore new revenue streams. We had a conversation with Raffaella Camera, Strategic Advisor Gaming & Lifestyle, former Head of Brands Epic Games' Unreal Engine, and Global Head of Strategy Accenture XR, and this section is based on what we learned from her.

The Expanding Role of Immersive Technologies in Retail

As consumer expectations continue to rise in the digital age, retailers are increasingly turning to Spatial Computing to create more engaging, personalized, and interactive shopping experiences. These technologies are no longer viewed as mere novelties; they are becoming essential components of a comprehensive retail strategy, driving innovation across the industry. However, Camera says that “at the end of the day, you have to show ROI. You have to be able to hit KPIs, and you need to be able to explain why you are spending marketing dollars there.”

Enhancing the Customer Experience—A New Era of Interaction

One of the most significant impacts of Spatial Computing is on the customer experience, where these technologies are bridging the gap between digital and in-store shopping. This convergence offers consumers a more immersive and personalized experience, fundamentally changing how they interact with products and brands.

Virtual Try-Ons—Bringing the Fitting Room to the Customer

Virtual try-ons have emerged as a key application of AR technology, allowing customers to visualize how products such as clothing, makeup, and accessories will look on them without the need to physically try them on. This technology not only enhances the shopping experience by providing a higher level of convenience and personalization but also aims at addressing common pain points such as fit and style, ultimately reducing return rates—a significant cost factor for retailers. As AR technology continues to advance, the accuracy and realism of virtual try-ons are expected to improve, making them an indispensable tool in the retail arsenal.

The founder of SyncReality, Cyril Tuschi, came to our house to show us how this could work. He scanned our home on his iPhone, then ran a virtual fashion show where virtual models were walking on our real floor right in front of us⁴. We could stop the show, grab the clothes off of one of the models, and buy them. Or try them. He sees a world where a digital twin of you can wear the clothes so you can see how they look on you instead of a virtual model that probably has a different body shape than you do.

Personal Avatars—The Future of Online Shopping

Taking virtual try-ons a step further, the development of hyper-realistic personal avatars is set to revolutionize online shopping. Camera says, “On the marketing side, personal avatars are becoming increasingly popular, and I believe it will soon be the norm to have multiple avatars per person, each for different purposes. These could range from seeing exactly how clothes fit or how your makeup looks on you, to more entertainment-based avatars you might use in social games or similar activities.”

Camera says, “These avatars, which accurately replicate a customer’s body measurements, skin tone, and facial features, will allow for even more precise and personalized virtual try-ons. Imagine being able to see exactly how a garment will fit and move on your body, or how a particular shade of lipstick will look against your skin tone—all within a virtual environment. This level of customization and realism could significantly boost consumer confidence in online purchases, leading to higher conversion rates and fewer returns.”

Virtual Tours and Room Builders—Redefining Real Estate and Home Decor

Camera says that beyond fashion and beauty, Spatial Computing is also transforming the way consumers shop for furniture, real estate, and home decor. Virtual tours and room builders allow customers to explore properties and visualize how furniture and decor items will look in their homes before making a purchase. This technology is particularly valuable in the luxury market, where high-value purchases often require a greater level of assurance and personalization. A prime example of this is the room-scale VR tour created by Four Seasons to showcase and sell its Private Residences at Lake Austin even before construction began. By providing customers with a realistic and interactive preview of their purchases, retailers can enhance the buying experience and build stronger customer relationships.

Operational Efficiency and Innovation—Streamlining Processes with Immersive Technologies

While enhancing the customer experience is a primary focus, Spatial Computing also offers significant benefits in terms of operational efficiency and innovation. These technologies are being used internally by retailers to streamline various processes, from product design and development to merchandising and marketing.

Design and Prototyping—Bringing Concepts to Life in 3D

One of the most powerful applications of Spatial Computing in retail is in the design and prototyping phase. By leveraging 3D modeling and simulation, retailers can bring product concepts to life in a virtual environment, allowing for rapid iteration and testing. This not only accelerates the design process and facilitates remote collaboration but also enables more accurate validation of materials, fit, and functionality. For example, automotive companies can use VR to simulate and test new vehicle components, ensuring they meet stringent safety and performance standards before physical prototypes are produced.

Virtual Showrooms and B2B Virtual Samples—Revolutionizing the Wholesale Process

The rise of virtual showrooms and B2B virtual samples has been a game-changer for retailers, according to Camera, particularly during the COVID-19 pandemic when physical interactions were severely restricted. These virtual solutions allow retailers to showcase entire product lines to buyers and partners in a digital environment, enabling real-time feedback and collaboration. This approach not only saves time and resources by eliminating the need for physical samples but also offers greater flexibility in terms of customization and iteration. As a result, virtual showrooms are becoming an increasingly popular tool for retailers looking to enhance their wholesale process and reach a global audience.

Merchandising and Store Layout—Optimizing Retail Spaces with Spatial Computing

Spatial Computing technologies are also being used to optimize merchandising and store layout, helping retailers create more efficient and effective retail environments. By using game engines, AR, and VR to simulate different store layouts and product placements, retailers can test various configurations and determine the most effective ways to display products and guide customer flow. This approach not only improves the overall shopping experience but also maximizes sales by strategically placing high-margin items in key areas of the store. Moreover, Spatial Computing allows retailers to conduct competitive analysis by comparing their store layouts with those of competitors, ensuring they stay ahead in the market.

Marketing and Brand Engagement—Creating Immersive Experiences that Drive Sales

In addition to operational efficiencies, Spatial Computing is revolutionizing retail marketing strategies, offering new ways to engage customers and drive sales. From virtual product launches to branded virtual worlds, these technologies are enabling retailers to create immersive experiences that resonate with consumers and build brand loyalty.

Virtual Product Launches—Extending the Reach of Physical Events

Camera says that virtual product launches have become a popular marketing tool for retailers, allowing them to extend the reach and impact of physical events. For example, luxury brands like Louis Vuitton have recreated their fashion shows in virtual environments, complete with digital replicas of their garments and accessories. These virtual events not only allow a global audience to experience the brand's latest collections but will also offer opportunities for interactive engagement, such as virtual try-ons and personalized product recommendations. By combining the exclusivity of a physical event with the accessibility of a virtual one, retailers can create a truly unique and memorable brand experience, extending its lifespan and geographical reach.

Branded Virtual Worlds and In-Game Advertising—Reaching Consumers in Their Favorite Digital Spaces

As consumers spend more time in virtual environments, such as gaming platforms and traditional social media platforms, retailers are increasingly exploring ways to reach them in these spaces. Branded virtual worlds and in-game advertising offer a powerful way to connect with consumers in a context that feels natural and engaging. For example, Camera says, beauty retailer Sephora has launched Sephora Universe, a branded immersive online destination that offers a seamless blend of innovation and learning where customers can explore new products and participate in virtual events. Similarly, brands are placing in-game advertisements and mini-experiences within popular games like Roblox and Fortnite, allowing them to reach a highly engaged audience in a way that feels organic and unobtrusive.

A Super League study shows that pop-ups in Roblox—mini-experiences within a published game—produce 14% engagement, three times more than TikTok.

Challenges in Implementation—Navigating the Complexities of Spatial Computing

While the potential of Spatial Computing in retail is vast, the implementation of these technologies is not without its challenges. When relying strictly on AR and VR technologies, retailers must navigate a complex landscape of technological limitations, consumer adoption barriers, and ethical considerations to successfully integrate these solutions into their business strategies.

Technological Limitations—Overcoming the Hurdles of Device Usability and Content Availability

According to Camera, one of the primary challenges in implementing Spatial Computing technologies in retail on AR and VR devices is the current state of the hardware and software. While these technologies have made significant strides in recent years, they are still not as user-friendly or widely adopted as traditional digital platforms. For example, AR glasses and VR headsets often have limited fields of view, short battery life, and cumbersome designs, making them less appealing for everyday use. Additionally, there is a lack of compelling content and practical use cases that would encourage consumers to adopt these technologies on a regular basis. Retailers must therefore carefully consider the cost and ROI of investing in AR and VR solutions, particularly when the consumer base for these devices is still relatively small.

Consumer Adoption Barriers—Building Trust and Confidence in New Technologies

In addition to technological limitations, there are also significant consumer adoption barriers that retailers must address. While early adopters and tech enthusiasts may be eager to embrace AR and VR, the broader consumer base might be more hesitant—especially when it comes to making purchases based on virtual experiences, unlike in gaming environments where users are already accustomed to buying digital goods.

Building trust and confidence in these technologies is therefore crucial for retailers. This can be achieved through education and transparency, as well as by offering guarantees such as easy returns and exchanges for products purchased through AR or VR platforms. By addressing consumer concerns and demonstrating the value of these technologies, retailers can encourage wider adoption and ultimately drive sales.

Ethical Considerations—Balancing Innovation with Privacy and Security

As with any emerging technology, says Camera, the adoption of Spatial Computing in retail raises important ethical considerations, particularly around privacy and data security.

These technologies often require the collection and processing of large amounts of personal data, such as body measurements, facial features, and shopping preferences. Retailers must therefore take steps to ensure that this data is handled securely and transparently, with clear policies on how it is used and shared. Additionally, as AI becomes more integrated with AR and VR, there are concerns about the potential for bias and discrimination in personalized recommendations and advertising. Retailers must therefore be vigilant in ensuring that their use of these technologies is fair and equitable, and that they do not inadvertently reinforce harmful stereotypes or practices.

The Future of Retail with Spatial Computing—Opportunities and Trends

Looking ahead, the future of retail will be increasingly shaped by the continued evolution of Spatial Computing technologies. As these technologies become more sophisticated and accessible, they will unlock new opportunities for retailers to innovate and differentiate themselves in a competitive market.

Personal Avatars and Hyper-Realistic Try-Ons—Redefining Personalization

One of the most exciting developments on the horizon is the rise of personal avatars and hyper-realistic virtual try-ons. As discussed earlier, these avatars will allow consumers to see exactly how products will look and fit on their own bodies, leading to a more personalized and confident shopping experience. In addition to fashion and beauty, this technology could be applied to other industries such as fitness and health, where consumers could use avatars to try out exercise equipment, test fitness apparel, or even visualize the impact of dietary changes on their bodies. By offering this level of personalization, retailers can create a more meaningful and engaging shopping experience that resonates with consumers on a deeper level.

Virtual-to-Physical Commerce—Bridging the Gap Between Digital and Real-World Shopping

According to Camera, another emerging trend is the integration of virtual-to-physical commerce, where consumers can seamlessly transition from a virtual shopping experience to a physical purchase. Retail giants like Amazon and Walmart are already exploring this concept, with initiatives such as Amazon Anywhere and Walmart's powered virtual stores in platforms like Roblox. In this model, consumers can naturally interact with products while playing a game, exploring a world, or trying products on their avatars, and then effortlessly purchase the physical items, which are delivered to their doorstep. This approach not only enhances convenience but also allows retailers to gather precious data on consumer preferences and tap into the growing demand for seamless and integrated shopping experiences.

As this trend continues to evolve, we can expect to see more retailers adopting virtual-to-physical commerce as a core component of their business strategies.

Cross-Platform Standardization and Regulation—Navigating the Regulatory Landscape

As the use of Spatial Computing in retail expands, there will be an increasing need for cross-platform standardization and regulation. Currently, the landscape for these technologies is fragmented, with different platforms offering varying levels of support and functionality. To create a more cohesive and interoperable ecosystem, says Camera, retailers, developers, and tech companies will need to collaborate on establishing common standards for content creation, data security, and consumer protection. Additionally, government regulations, particularly around data privacy and content moderation, will play a crucial role in shaping the future of these technologies. As regulatory frameworks evolve, retailers will need to stay informed and compliant to ensure the responsible and ethical use of Spatial Computing in their operations.

The Shift Toward Experience Platforms—Creating Immersive Ad Ecosystems for Brands

Finally, we may see the rise of experience platforms, as Camera defines them—a new breed of immersive ad platforms that allow brands to integrate their products across various popular games already published in platforms like Roblox and Fortnite. Camera says that these platforms could offer a range of experiences, from mini-games and branded virtual worlds to in-game advertisements and interactive product placements. For example, a fashion brand could create a virtual boutique as a popup in a popular game, or sponsor an exclusive virtual event, insert branded NPCs (nonplayer characters), or branded visual effects, and interact with other brand enthusiasts. By leveraging the success and audience of an already established game and standardized immersive ad units, retailers can reach consumers in new and engaging ways, minimizing their risk and investment.

The Extension of Branded IP in Virtual Environments

The demand for 3D-based products in commerce is set to rise. Brands should identify which elements of their product lines and IP can be transformed into virtual goods or accessed virtually. For example, the automotive industry can incorporate 3D entertainment and tools into vehicles, emphasizing comfort, connectivity, and entertainment. Fashion can further explore virtual clothing and virtual makeup, while the music industry can offer virtual tours, virtual merchandise, and emotes. A notable example is LEGO, which extended the value of its IP by creating and monetizing a game in Fortnite, offering LEGO templates and props to other game creators, and selling LEGO Fortnite sets in-store.

Autonomous Cars and the Future of Retail—A New Frontier for Spatial Computing

As autonomous vehicles become more prevalent, they present a unique opportunity for retailers to rethink the shopping experience. Autonomous cars are essentially mobile living spaces, offering passengers time to engage in various activities, including shopping. Retailers could explore ways to integrate their offerings into the autonomous vehicle experience, allowing passengers to browse products, make purchases, and even have items delivered to their destination. For example, Camera says, an autonomous car could be equipped with AR displays that allow passengers to view and interact with products from a specific retailer, or even place orders for pickup at their next stop. This concept, while still in its early stages, represents an exciting frontier for Spatial Computing and retail innovation.

Embracing the Future of Retail

Spatial Computing in retail is not just a fleeting trend; it represents a fundamental shift in how businesses engage with customers, manage their operations, and expand their market reach. While there are challenges to overcome, the potential benefits—ranging from enhanced customer experiences and operational efficiencies to new revenue streams and innovative marketing strategies—make these technologies a worthwhile investment for forward-thinking retailers.

As we move into the future, the line between the physical and digital worlds will continue to blur, and retailers who embrace this shift will be well-positioned to thrive in the new era of commerce. By staying ahead of technological advancements, navigating ethical considerations, and exploring new opportunities for engagement, retailers can create a more dynamic, personalized, and immersive shopping experience that resonates with the modern consumer.

Spatial Computing for Logistics

Logistics is undergoing its own emerging mini-revolution when it comes to utilizing Spatial Computing. Companies like Seegrid, who make **Vision-Guided Vehicles (VGVs)**, help in reducing the high accident rates associated with human-manned forklifts. These VGVs use Computer Vision and sensing technology to know where to go on the floor and to avoid obstacles in their way. Companies such as BMW, Jaguar, and Whirlpool, among many others, use Seegrid's VGVs on their warehouse floors.



Figure 7.4: Photo credit: Seegrid. Examples of Seegrid's VGVs that replace human-manned forklifts.

Another Spatial Computing warehouse helper is the GreyOrange Butler robot, which GreyOrange calls a “Decision Science-Driven Robotic Goods-to-Person System” for “automated put-away, inventory storage, replenishment, and order picking in fulfillment and distribution centers.” GreyOrange Butler customers are mostly retailers with physical and/or online stores, in addition to apparel makers, third-party logistics providers, and consumer electronics companies. Potential competition for both Seegrid and GreyOrange logistics products comes from start-ups like Zebra Technologies and 6 River Systems, as well as larger established players such as Swisslog.



Figure 7.5: Photo credit: GreyOrange. GreyOrange Butler robots working inside a warehouse.

In terms of logistics' use of AR headsets, there has been some version of AR used in warehouses since at least 2008, though the overlay visuals used were in 2D versus 3D up until the last few years. Companies such as RealWear and Microsoft, with the HoloLens 2 headset, along with the now-defunct companies ODG and DAQRI, have catered to this market that we predict will have exponential growth within the next five years due to a fast-growing need for logistical warehousing systems.

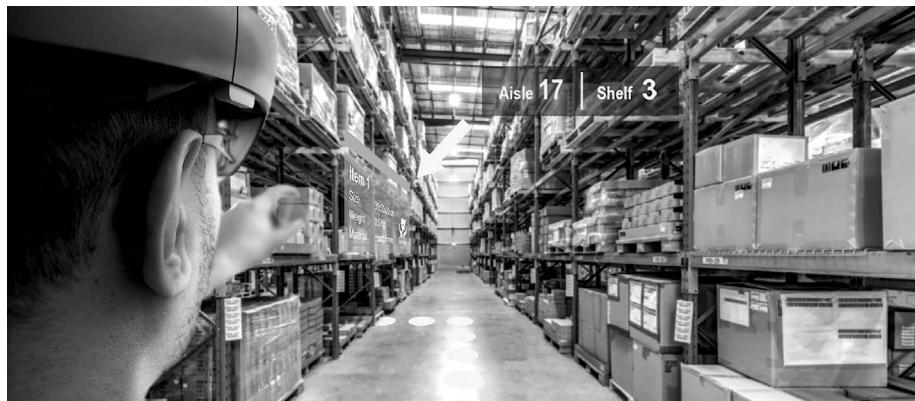


Figure 7.6: Photo credit: ARM23. Example of an AR program for logistics using a Microsoft HoloLens headset.

Underlying this demand is an ever-increasing growth in online sales for a wide diversity of products.

Spatial Computing Delivery Methods—Robots, Drones, Future

Starship Technologies is a venture that was started in 2014 by Skype co-founders Janus Friis and Ahti Heinla, to address the “last mile” of delivering goods such as groceries, restaurant take-out, and packages. Their product is an autonomous electric robot that navigates city streets and paths using six wheels, Computer Vision, and a sensor suite that includes cameras, GPS, and an inertial measurement unit. It does not use LIDAR in a bid to cut down on costs to the customer.

We talked with Sean Eckard, former Senior Manager of Business Development at Starship Technologies. He said it is “supreme in every way” to deliver things with a 50-lb electric robot instead of a 2,000-lb car.

Funny enough, it isn't the technology we mostly talked about while we were with him. Rather, it is the regulatory approach that he's seeing from different cities. Some cities met his company with open arms, recognizing the benefits that delivery robots will bring to their citizens: lower costs and more consistent delivery times, along with new jobs that come from companies like Starship. Others, however, refused to consider letting Starship's robots roll down their streets, seeing them as too dystopian, a threat to existing jobs, or a potential safety hazard.

Eckard said the job is harder for some of Starship's competitors since they made their robots bigger and uglier. Starship has overcome some of these objections, Eckard said, because of its fun low profile and attention to detail when it comes to pedestrian safety, which is mostly due to a ton of AI training on what to do around, say, kids or dogs. He says that is why they started on college campuses before going to entire cities. They wanted to have the ability to watch and train the robots in a constrained space, which also gained the trust of people using them.



Figure 7.7: Photo credit: Starship Technologies.

Oh, and by going to college campuses, they assumed they would have a lot more trouble with thefts and abuse, which turned out not to be true. In fact, they are having trouble getting anyone at all to rip them off (and were actually hoping for more so that they could train the robots on what to do if someone carted one off in, say, a pickup truck). It turns out that even most criminals are smart enough to know that the cameras onboard are recording them, just the same as a new Tesla records someone breaking a window or doing other damage to the car.

It was Eckard's opening argument, though, that's going to be most interesting to watch: why are we driving cars around to go pick things up, or deliver them, when a 55-lb robot can do it just as well? One, presumably, that will take a lot less energy to move around, which means less carbon put in the air for each package delivered.

Climate change is heating up as an issue, as it were. This increased pressure to do something about climate change will lead more companies to use these delivery robots, it seems to us, and make other changes as well to reduce the amount of carbon that they are putting in the air as they move products to people's homes.



Figure 7.8: Photo credit: Amazon. An Amazon Scout robot reported to have been built on the acquired IP of a company called Dispatch.

Amazon announced Amazon Scout, a very similar-looking autonomous robot (also called a self-driving delivery device, or just a robot) in January 2019, and was running small field tests with it on getting packages to customers in Snohomish County, Washington, and Irvine, California. It predominantly used LIDAR as its remote sensing method, which Starship feels is overkill. Amazon cancelled Amazon Scout in January 2023 due to cost-cutting measures.

Some are definitely looking at delivery robots as a way to completely rethink business.

Take, for example, pizza.

Alex and Julia Garden noticed that the average delivery took about 20 minutes; about the same length of time it takes to cook. In a bid to completely disrupt the pizza supply chain, they outfitted their Zume delivery vans with dozens of robot-controlled ovens. The pizza is cooking as it drives to your house. This reduces the capital, since they don't need an expensive kitchen and building to house it, and their pizzas arrive hotter than their competitors', literally coming out of the oven at your front door. Zume, alas, will no longer be doing these kinds of operations due to internal mismanagement.

We can imagine they won't be the last business to completely rethink retail, delivery, or both thanks to Spatial Computing.



Figure 7.9: Photo credit: FedEx. An example of a FedEx last-mile delivery robot, Roxo, the SameDay Bot.

Roxo, FedEx's SameDay Bot, was its answer to the last-mile delivery issue. FedEx had been working with AutoZone, Walgreens, Target, Pizza Hut, Lowe's, and Walmart to adapt automated delivery technology to customers' needs. What made Roxo a little different than the other delivery robots is that it can travel upstairs to arrive directly in front of a person's door. FedEx stopped the development of Roxo in October 2022, also due to cost-cutting measures.

Amazon's vision of 30-minute or less package delivery is encapsulated in the design of its drones. Its drone operation is run by Amazon Prime Air, a subsidiary of Amazon. It currently offers drone services in College Station, Texas, and Tolleson, Arizona, with plans to start in a third U.S. city and the U.K. and Italy by the end of 2024. Our friends who have gotten a look at the Spatial Computing technology inside say that what's really fantastic is the paper-thin radar sensor inside that maps out the world in 3D and can fly to your door the same way your AR glasses will soon bring all sorts of virtual fun to you: through new technologies such as "AR Cloud" and **Simultaneous Location And Mapping (SLAM)**. Already, consumer drones are starting to use the same techniques to navigate through trees or around obstacles.



Figure 7.10: Photo credit: Amazon. An example from Amazon Prime Air: a future delivery system designed to get packages to customers in 30 minutes or less using unmanned aerial vehicles, also called drones.

In April 2019, Wing Aviation, the drone delivery arm of Google's parent company, Alphabet, became the first American company to be granted FAA approval for commercial delivery. In addition to Amazon, Uber and UPS have applied for air carrier certificates from the U.S. Federal Aviation Administration to launch commercial drone operations.

Prior to having deliveries of goods come to you, Spatial Computing experiences could be used to visualize what furniture and other goods would look like in your living room and anywhere else.

In Your Living Room and Elsewhere—Spatial Computing Marketing Experiences

We *wrote* this part of the chapter by talking into a microphone to produce text (elsewhere, we generally just used a keyboard). This is to get ready for the world that Amazon and others are bringing to us. Alexa listens to us and brings us whatever we want. “Hey Alexa, put a box of Cheerios on my shopping list.” “Hey Alexa, can you deliver some Chinese food?” “Hey Alexa, can we watch *Halt and Catch Fire*?” “Hey Alexa, I want to buy some clothes.”

Starting in the Living Room

Ubiquitous computing is already changing our living rooms. Robert Scoble says, about his living room: “I have a Google device listening to me. I have many Apple devices listening to me. I have a TV that listens to me. I have a Sonos soundbar and a Facebook Portal device that listens to me and, of course, we have Amazon boxes listening.”

Soon, we will have three-dimensional sensors on our faces and in our living rooms and new kinds of displays on our eyes, potentially replacing our television sets and computer displays. Hooray! There’s a company called Looking Glass that has a flat display that gives you a three-dimensional look at three-dimensional objects. And other devices or technologies have been shown to us that include three-dimensional contacts that don’t need to be on your face. But it is the on-the-face wearable devices that are the most interesting. They are also the most controversial, because of the privacy implications of having a three-dimensional sensor that can literally see everything in your house. If you take it to the utmost, there are wearable devices that are going to remember where you left your keys. Or, where you left your Pepsi, or how many bites of food you’ve taken in the last hour.

They could also know who has been in your house or how many times you’ve actually interacted with your children or your spouse. What music you listen to. What things you bought. What tasks you put on your task list. What things are on your calendar. What you’re wearing on a daily basis. Or, if you’re wearing anything at all. They might catch wind of your emotions. Did you just yell at your children? Do you feel guilty about it? Did that cause you to buy some Ben & Jerry’s ice cream? Soon, Taylor Swift might start playing in your ears singing “You Need to Calm Down” and ice cream ads might pop up with messages such as “need a break?” because you are raising your voice.

Already, we can use a Meta Quest 3 or Apple Vision Pro to watch movies on Netflix, or TV shows, music videos, or other kinds of content on our personal devices.

The children can watch SpongeBob while the adults may watch something a little more racy. What does this mean for personal interaction? And what does it mean for retail?

A Personalized Experience

We are heading into a hyper-personalized world where Amazon, Google, Facebook, Apple, and others will know a ton about our lifestyles, our family structure, and our favorite media choices. Adding to that, Amazon already knows everything you've purchased from Amazon and, now, its member brands, such as Whole Foods.

As Jeff Wilke, Amazon's former CEO of Worldwide Consumer, said, "We're going to see AI and Machine Learning embedded in all kinds of different products, so it makes personalization better. We just launched an enhancement to the iOS app that lets you screenshot a fashion look that you're interested in and find something on Amazon that matches."

So, let's go back to the three-dimensional store that will soon be shown to us in our living room. It can predict what you're going to buy next with some level of confidence. Planning a trip to Rome? It will suggest new luggage. Running low on groceries? It might suggest you look around a new kind of Amazon store where you could pick out strawberries, Cheerios, and your favorite coffee. All to be delivered in an hour or less by a robot.

Many people will find this very attractive because heading to the store often means being stuck in traffic in cities or stuck in lines at grocery stores and this opportunity cost of time ends up being quite high indeed.

Of course, as you use more of these services, they learn more about not just your shopping behavior, but everyone's shopping behavior. This is already how Amazon works and why Amazon is already beating the experience of shopping at a real store like Fry's.

These new three-dimensional stores that are going to be in your living room are also going to mean that marketers are going to have to think differently. Soon, we will be able to have conversations with virtual beings: things that look like real human beings and can actually have a pretty good conversation with you because they know everything about what you're buying. Your virtual being could ask the same questions I asked thousands of people in a retail store in the 1980s. "Would you like a small pocket camera? Or a big DSLR?" "Would you like to see a Nikon or a Canon?" "Do you know the difference?"

It could warn you that some of your product choices are not in line with your new intentionality. You could tell it, “Hey, I want to lose 10 pounds within the next month,” and it could tell you that it’s not a good idea to buy that Ben & Jerry’s ice cream that it “knows” you want to buy.

Facebook and other big companies are already planning for such a world years in advance. We’ve caught wind of design documents for this kind of future world. The designers inside these companies are already thinking about how your buying behavior will change thanks to these new technologies in your living room.

Steve Conine, co-chairman and co-founder of the online retailer Wayfair, says, “My belief is, at some point in the not-too-distant future, every home in America will have a three-dimensional model tied to it. If you give us access to it, we can use that to show you styles and designs in that space that look amazingly good.”

That’s for the near future and the farther-out future. Let’s now go back to what is currently happening with regard to Spatial Computing. In this case, we will focus on how some major select furniture and accessories stores are using AR, and, in some cases, VR.

The VR World

In 2016, IKEA came out with its Virtual Reality Kitchen Experience, where potential customers could browse the IKEA catalog to place items around the virtual kitchen, open drawers and oven doors, fry and toss around meatballs, and make pancakes. The app was developed for the HTC Vive and released on Valve’s Steam platform. The experience was not widely known about, however, as an article published in the Financial Times in April 2018 made it appear that this experience was a new one.

In December 2017, for the opening of its newest store near Dallas, IKEA partnered with media and technology agency Wavemaker to offer VR games and experiences to potential customers entering the store. Close to 300 people played a “pillow toss” game or hung out with a virtual panda. With this, IKEA was able to test what types of VR content would attract potential customers.

This is part of a larger effort to test and learn what kinds of investments should be made with regard to IKEA's stores, logistics, digital capabilities, and services.

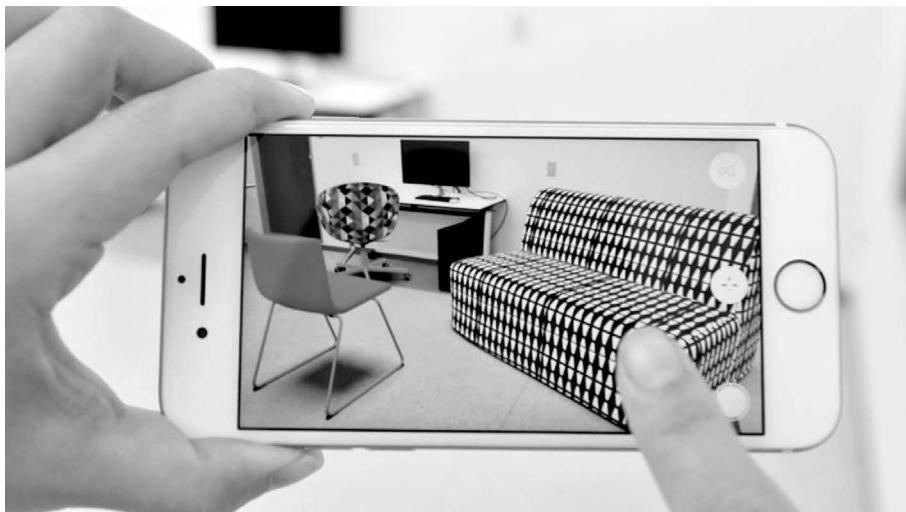


Figure 7.11: Photo credit: IKEA. Advertising for IKEA's "Place" app.

With regard to its mobile AR app, IKEA has collaborated with many partners, including ad agency 72andSunny, to help prospective customers better visualize IKEA furniture in their homes and to foster strong consumer loyalty.

In September 2017, IKEA launched its AR app, IKEA Place (now IKEA Kreativ) for Apple iOS devices. Over 2,000 products—nearly the company's full collection of sofas, armchairs, coffee tables, and storage units—were available to view and place. In contrast to the catalog AR app IKEA launched in 2013, the placements were more stable, staying put on the floor, and in their proper measured proportions. Also, in 2013, there were only about 100 products versus the more than 2,000 products currently available to view and place. Also, the products appear with about 98 percent accuracy in scale and true-to-life representations of the texture, fabric, lighting, and shadow.

The number of IKEA products in IKEA Kreativ available to view is more than 3,200, with the AR app also available for Android mobile phones. IKEA's VR strategy has not been as clear as its AR strategy. Wisconsin's first IKEA opened on May 16, 2018, in Oak Creek, Milwaukee. As part of the store's opening festivities, visitors to a pop-up event close by in McKinley Park could participate in a VR experience that displayed an IKEA showroom and designs, as well as some of its latest products.

As the world's largest retailer, Walmart is also a company competing on multiple fronts. Its culture is nearly the opposite of Amazon's, having begun in the U.S. South, and building significantly on a foundation of price-sensitive rural shoppers.



Figure 7.12: Photo credit: Walmart. Two screenshots from the original Spatialand (now spelled "Spatial&") app.

With this app, consumers could try outdoor products in a Yosemite National Park VR experience. Yet Walmart transitioned well into the online space, where its website is the third largest after Amazon and Apple, better than all of its traditional shoppers in terms of visitors and sales. Its e-commerce business, which accounts for 3.6 percent of its U.S. sales, has incurred much expense to grow. Its website has about 100 million unique monthly visitors, versus 180 million for Amazon.com. However, recent reports indicate that despite significant and continuous online growth, investors remain unsatisfied that it can stop Amazon's relentless rise.

Walmart keeps trying and expanding, showing it is unafraid to take risks with new technologies, particularly AR and VR. In 2017, it acquired Jet.com, an online shopping network often compared favorably to Amazon. Marc Lore built Jet.com and came on board as head of Walmart's U.S. e-commerce business in September 2017 following Jet.com's acquisition. In the third quarter of 2017, Walmart's U.S. e-commerce sales grew 50 percent year on year.

Making an Impact

To devise new technologies that have an impact on the way people shop and thus hopefully increase sales, Lore created Walmart's in-house start-up incubator called Store No. 8, named for the store that founder Sam Walton had once used for experimentation.

Much focus had been on VR solutions. However, in January, 2024, Store No. 8 was closed down due to cost-cutting. VR within Walmart has been replaced with investigating AR initiatives. With regard to Amazon, its strategy and culture have always been to use technology to become more efficient and improve customer experiences. It has been lethal to competitors on virtually every consumer-facing front. As always, Amazon moves successfully in multiple directions: it has expanded successfully into B2B with AWS, a wholly-owned web services subsidiary. It is building brick-and-mortar stores and has acquired Whole Foods, the supermarket chain; now it is buying large goods producers to lower grocery costs for consumers. On another front, according to job postings, Amazon is creating statistical three-dimensional models of human bodies, which it will then match to images and videos of people via Deep Learning algorithms and other tactics. That will have a “wide range of commercial applications” for Amazon customers, the company says in the job ad.

Among the company’s many competitive assets is that it probably has more data on what people buy than any company on Earth, and with the addition of Amazon Echo, it now has a friendly device that people talk to—and in doing so, provide massive amounts of additional data on a growing number of the world’s households.

Like all major IKEA competitors, Amazon has an AR application called “AR View” for both iOS and Android; we tried it for a few minutes on an Android phone and found it to be pretty similar to IKEA Place. Amazon’s many thousands of items to view are welcoming. If the current app was the end-game for Amazon in AR and VR, IKEA could comfortably remain the leader of a small, but aggressive, pack of competitors. But Amazon’s historic long-view strategy would make that possibility unlikely. Where they are headed, we can only guess, but we speculate that a future iteration of AR View will include headset applications. Even more significant is that it will eventually link Echo to AR View. It does not appear that Amazon will be utilizing VR for retail purposes, with earlier indications pointing to its interest there solely for the entertainment industry.

In a few years, pattern recognition will be able to predict what consumers will need next, whether that is another quart of milk or a new couch. Alexa could actually suggest that it is time, and suggest they look at AR possibilities of an 84-inch couch that would fit properly in a certain space, then show colors that would match the curtains behind it. This could all happen before a family even considers shopping for the couch online or in a store. While competing with Amazon on this level, we also see an opportunity for partnering elsewhere via AWS, Amazon’s B2B web services subsidiary.



Figure 7.13: Photo credit: Irena Cronin. A chair sold on Amazon being viewed through the company's AR app, "AR View" (iOS version).

AWS' Amazon Sumerian is a non-technical platform for creating AR or VR applications without programming. Imagine that we have a room hosted by an AI-enriched host avatar. The rooms are easily filled with three-dimensional objects and, from what we can discern, the room can be as large or small as developers wish to make it. Perhaps this could be a furniture showroom. The host can be a sales representative who shows online shoppers selections from any chosen category. Sometime in the near future, haptic technologies could be added on, perhaps in the form of a glove. Then, shoppers could actually feel textures and surfaces.

A “No Install Needed” AR World

In late 2016, Amazon began planning for what AR and VR should look like for AWS customers, with a preview debuting at AWS re:Invent. They then bought Goo Technologies, a bankrupt Swedish company. Goo Create, its three-dimensional creation platform, became the foundation of Sumerian's **Integrated Development Environment (IDE)**. One of the biggest benefits of a scalable cloud infrastructure is to be able to dramatically reduce latency, a feature that is very important for quality presentation and manipulation when using mobile AR apps. Another huge benefit is that you won't need to load up a new app just to, say, see some functionality in a shopping mall.

With regard to open standards, Sumerian supports several of them: WebGL, WebAR, WebVR, and the emerging WebXR framework that brings AR/VR apps universally to all devices and browsers across platforms.

In many ways, Wayfair Furniture is the cultural opposite of Walmart: it prides itself on having its engineering team at the core of its organization and credits engineers for “creating customer experiences that we would enjoy ourselves.”

Founded in 2002, it fits the classic start-up model. Two high school pals go off to college together and fall in love with technology. They start on a shoestring budget with a single home furnishings boutique, then another, and another. They consolidate into a single online home-furnishings superstore, go public, and become one of the most successful retailers in modern history.

Today, they offer over 10 million products from over 10,000 suppliers, and yet the two founders who continue to head the company see Wayfair as a tech company, rather than a home furnishing merchant.

They are sometimes credited with inventing the term “v-commerce” to replace the older “e-commerce” term because the concept of a virtual store for real shoppers is what they say they are all about. We point to this because we believe it gives insight into their corporate culture and priorities.

Accordingly, the company has been early to introduce AR over VR apps. The “View in Room 3D” app was launched with ARKit at the same time as the IKEA Place (now Kreativ) app. While it did not receive the coverage and favorable reviews that IKEA Place received, it was generally perceived as a quality offering with a broad selection of three-dimensional rendered items for sale. As with other competitors included in this report, Wayfair introduced an Android version of its AR app based on ARCore in early 2018, following its iOS ARKit version introduced in the fourth quarter of 2017.

Now, Wayfair has a new app called “Decorify,” which uses Generative AI and Computer Vision to place furniture and other items in your room. Launched in July 2023, and initially for mobile and desktop via browser use, it is now available for the Apple Vision Pro.

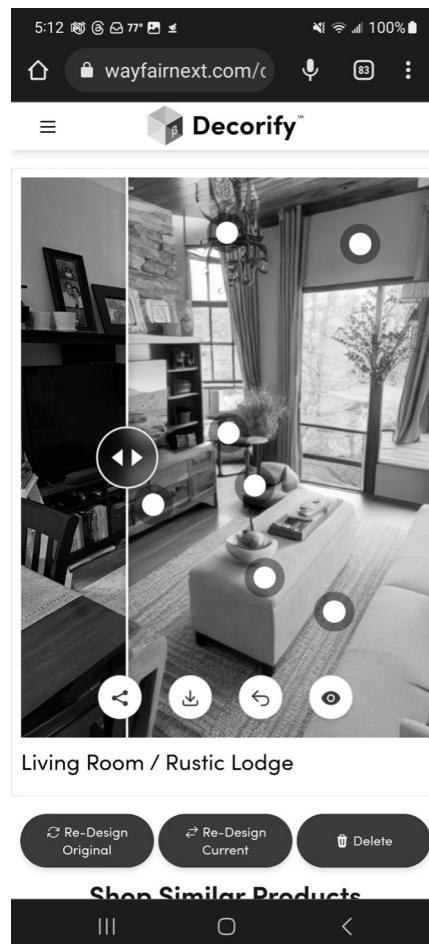


Figure 7.14: Photo credit: Wayfair. Screenshot of the Decorify mobile app.

Wayfair has also been quick and aggressive on all popular social media platforms, where 71 percent of all consumers say they are influenced on buying decisions. Understanding that visual product representations are extremely important to marketing and sales, it has made Instagram a top priority and achieved its greatest growth on any social platform there.

Founded in 2010, by a Silicon Valley couple who became frustrated with trying to find useful information online to help them with home renovation, Houzz is a community of home improvement people, architects, contractors, furniture, and home accessory providers.

It started on the iPad and did not get to Android until 2013. By then, the community had helped more than six million people. Houzz is focused on helping buyers and makes most of its money through commissions from selected vendors and from its online marketplace.



Figure 7.15: Photo credit: Houzz. Advertising for Houzz's "View in My Room 3D" app.

It began to go international in 2014. Today, it has localized versions for the UK, Australia, Germany, France, Russia, Japan, Italy, Spain, Sweden, Denmark, India, Ireland, and New Zealand.

Since its start, the online-only community platform has been prolific with innovative products, most of which are not currently on AR or VR, but have the potential to be used in that way moving forward.

These include the following:

- **Real Cost Finder:** An interactive tool that helps users plan home renovations and calculate total costs based on data collected from the Houzz community
- **Site Designer:** A free website-building and publishing tool for home professionals

- **HouzzTV:** This is essentially videos on room designs that we assume will be immersive in one form or another
- **Sketch:** A feature that lets users upload their own photos or choose any from the platform to add products from the Houzz Marketplace and see how they would work in rooms being redesigned
- **Visual Match:** A visual recognition tool that uses Deep Learning technology to analyze more than 11 million home photos on Houzz and drag and drop two-dimensional versions into the “View My Room 3D” app

Each of these items would clearly work better in AR than they do currently. Based on Houzz’s embrace of emergent technologies, we assume that development work is already going to upgrade one or all of these.

As is the case with every company mentioned so far, the Houzz “View in My Room 3D” AR app for iOS and Android does essentially what IKEA Place does, and from our testing of it, it works quite well, offering over 500,000 images rendered in 3D. Houzz’s focus is on AR, not VR, though an earlier version of “View in my Room 3D” was also available to be viewed in VR.

The retail Spatial Computing experiences discussed here are helpful when trying to make a decision as to whether or not to buy something. They tie into our view that in the future, physical retail stores will eventually be non-existent, except for, perhaps, those that sell high-priced luxury and/or unusual items.

The Past and The Future

Much of the future of retail has its roots in what exists today. People will continue to have a need or desire to buy something, whether it be a pear, a pair of pants, or a sailboat. It also seems that the need and urge to buy things in stores traverses time. Going back a few thousand years, archeologists have found remnants of stores that were set up in Ancient Greece and Rome. Currently, we still have physical stores where people go to buy things. Going back to the early 1900s, there were some mail-order catalogs where you could order items. The Sears, Roebuck and Company catalog was a very prominent example. Sears, Roebuck and Company had actually started in 1892 as a completely mail-order company, with its physical retail stores only opening up starting in the 1920s.

The idea of ordering something without having to go to a store and having it delivered directly to you is certainly not a new one.



Figure 7.16: Photo credit: Sears, Roebuck and Company. Early 1900's Sears, Roebuck and Company Household Catalog.

The clothing company J.Crew, from 1983 to 1989, existed only as a mail-order company. The modern incarnation of the mail-order company is one that places its full catalog online exclusively, such as Amazon. More commonly, hybrids exist where physical stores, physical catalogs, and internet sales all contribute to a company's revenues.

Commonly heard, though, are stories of companies who want to have a much more fully fledged and thriving digital presence above all, such as Nike and Walmart. This is just the start of the digital model fully eclipsing the physical model when it comes to retail operations.



Figure 7.17: Photo credit: Los Angeles Public Library. Pedestrians crowd the intersection of Seventh Street and Broadway, Downtown Los Angeles, circa 1937.

It used to be that the downtown areas of major cities were where people would congregate to go on regular buying sprees and have their fill of social interaction. Retail tasks would need to be planned for and lists were made so that nothing would be left out—each trip made was usually coordinated and well thought out.

The downtown areas of major cities thrived on this constant influx of retail buying power. Whole industries and businesses grew around it—movie theaters, restaurants, and circuses, to name a few.



Figure 7.18: Photo credit: Public domain. B. Altman and Company department store in New York City, circa 1948.

Department store buildings were made so large that they could accommodate the vast variety of different objects for sale, as well as the large number of prospective incoming shoppers.

Most of these large behemoths have ceased to exist, though Macy's Herald Square in New York City is still standing since 1901 (albeit some of the original space has been leased out to other companies).

It appears that these days, the large spaces for retail are its logistics spaces—non-public warehouses that hold items in a very organized and project-oriented way. Spatial Computing appears to be getting a major toehold not only in warehouse and packaging logistics but also along the full arc of retail fulfillment, which includes the sampling of goods via three-dimensional visuals followed by the ordering of them—all done digitally.



Figure 7.19: Photo credit: Amazon. An Amazon logistics center in Madrid.

The future of retail belongs to those companies that seamlessly and transparently offer Spatial Computing systems that allow for frictionless retail fulfillment. When goods are received faster, joy could ensue.



Figure 7.20: Photo credit: Apple. The glass window design of the Apple Park building that is prescient of the Spatial Computing device it released in 2024.

One company that is due to capitalize on this is clearly Apple. With its devoted loyal customer base and the Apple Vision Pro, an “army” of people using their Spatial Computing technology will be ready to purchase retail goods seamlessly, using eye-gaze selection, hand gestures, and voice commands. The days of going to a physical store are truly numbered.

Will Anything Remain the Same?

As we look at all the changes coming to retail in general, we see that Spatial Computing will bring major changes to the supply chains that feed goods to stores, to the stores themselves, to the experience of using the stores, either in physical locations or at home, and to the distribution networks of trucks and robots that will bring goods to our homes and offices. There are lots of changes ahead—ones that are keeping many an executive up at night. These changes might seem fairly mundane compared to what’s potentially coming to healthcare—all hitting in the 2020s. Will anything be the same again? “No” is our prediction, as you’ll discover in the next chapter.

References

1. Snap spectacles: https://www.youtube.com/watch?v=0d-_nZpmU9Q
2. Sephora Virtual Artist app: <https://www.youtube.com/watch?v=NFApcSocFDM>
3. Amazon Lens: <https://www.ispot.tv/ad/5ZhV/amazon-lens-find-similar-styles>
4. SyncReality: <https://www.youtube.com/watch?v=qUFygTvE1Jc>

8

Vision Five: Virtual Healthcare

As we imagine a world where every doctor, nurse, administrator, and patient wears Spatial Computing glasses, we see five changes: the medical system becomes more efficient, patient care gets better, those who provide care are better trained, the patient experience improves, and healthcare becomes more predictive and personalized.

Spatial Computing will eventually redefine what we think of as healthcare and the medical system that delivers it. Every part of it will be touched, from the surgery rooms to the doctors who save your life. They have been using the Microsoft HoloLens 2 and already are starting to use the Apple Vision Pro to prepare for or perform surgery.

Nurses will be assisted by robots bringing them drugs, needles, or other supplies without forcing them to walk down the hall to gather them. All of this will be directed by glasses flashing visuals about their job process and their patients into their eyes, as well as audio going into their ears providing information and commands.

Patients will see their waiting room experience transformed thanks to VR, not to mention the fact that they might not need to wait at all, thanks to new kinds of telemedicine that they will be able to get from home or work. What we think of as healthcare might turn into a new kind of virtual coach, that is, a virtual being that helps you make better choices in life, for example, directing workouts so you meet your fitness and health goals¹.

Administrators will finally see efficiency gains from the move to electronic health records and will see nice returns on their investment in systems to actually change patient behavior, getting them to exercise more and eat better, thanks to games on their glasses.

We will have a new understanding of healthcare, one that's more predictive and can anticipate your healthcare needs based upon your current lifestyle, all thanks to devices you wear, much like today's Fitbits and Apple Watches. This will bring us a world where healthcare is more readily available, much improved, and cheaper. This will be a particularly radical change for the world's poor, in contrast to their current experience of healthcare.

Although this pipe dream is probably a decade away, our discussions with doctors, administrators, and others demonstrate that R&D is already well underway to bring about a new kind of healthcare, which will bring us new techniques that early studies show is even changing our ideas of how pain will be dealt with. Someday soon, instead of being prescribed opioids or other types of prescription drugs to tamp down your pain, researchers see a new kind of practitioner prescribing that you use VR instead. This may seem an odd way to relieve pain, but we'll describe how promising results are already being observed for pain relief in the *Brain Tricks/Virtual Medicine* section of this chapter.

Even better, we are about to head into a world where the healthcare system isn't reactive like it is today. Today, we mostly go to a doctor because we aren't feeling well, or to get annual checkups, which mostly include blood workups, weight and blood pressure checks, and some other simple procedures to look over your body. Worse, hospitals primarily deal with patients who are having heart attacks or other health failures. Many of the researchers we talked to see a new, predictive healthcare that will see problems in advance of a major issue, or even work to get you to live in a more healthy way, which would save a lot of money and lead to better outcomes for everyone.

If your glasses say "hey, you need to get to a hospital now," that is a far better situation than having to be brought to a hospital by an ambulance later, which costs a lot of money and may already be too late. Even in a situation where you might need surgery, the doctors will have your heartbeat and other data from the past few weeks available, allowing them to see much more of your history than if they were to hook you up to monitoring equipment and guess what's been going on based on a far smaller amount of data.

With all of these highly significant applications for Spatial Computing in healthcare, we will now give you a sweeping overview of what changes we can expect to see in the coming years, based upon our research and discussions with medical professionals.

The Virtualist

Dr. Sherylle Calder in South Africa is, in our eyes, the prototype of a new kind of medical practitioner who works with Spatial Computing technologies to improve her patient's lives—one that Dr. Brennan Spiegel, director of health research at Cedars-Sinai, calls "the Virtualist."

Spiegel cofounded a conference that deals with the changes coming to healthcare thanks to Spatial Computing and is writing a book of his own where he details this new kind of specialty.

Dr. Calder is an eye doctor who has found a new way to improve her patients' lives by manipulating their eyes—all with digital patterns on screens. She is the first sports scientist to be awarded a doctorate in visual motor performance and works with professional athletes, and now everyday people, thanks to the company she started, EyeGym². The aim of her product is to improve her patients' eyes and the part of the brain that deals with the visual input they provide. She does her work with patients completely on digital screens and has had remarkable results. The owner of the South African cycling team told us that she took its worst rider, in terms of falls per race, and, after spending a few minutes a day with Calder's system, the rider went from the worst to the best within a few months. Her digital system fixed a flaw in his brain that caused him to fall during races. While not purely "spatial," we believe her system demonstrates how computing can be used to heal the brain, and we will see applications like this greatly expanded as more people wear computing on their eyes.

She explained that more than 30 percent of our brain has to do with visual processing and that she found ways to identify flaws in that system and "exercise them" with patterns and colors on screens. She's the first doctor we saw using a completely digital approach to medicine, which is why we see her work as being groundbreaking. Working with the 2003 England Rugby World Cup team, her system wrought huge changes. They ended up winning that year in the final against Australia. She also told us she worked with Ernie Els during a low period of his golf career, and her work helped him to win a major within six months.

Her story of finding ways to improve brains with digital screens was verified in our discussions with many doctors who are doing the same to find ways to provide relief to those who suffer from depression, pain, autism, dementia, PTSD, and other maladies, and we see her and her work as a prototype of this new kind of specialist.

She won't be alone—several doctors and medical researchers have told us. If you walk through a modern hospital, you will see dozens of specialists, from radiologists to pediatric surgeons. Dr. Spiegel says that to understand this new field that is in its beginnings, we have to see Spatial Computing R&D being done in surgery rooms, waiting rooms, burn recovery centers, mental health clinics, and other places in the healthcare system.



Figure 8.1: Photo credit: Robert Scoble. Dr. Sherylle Calder, left, in her South African office, explaining the work she does to improve professional athletes' perception systems.

For instance, UCSF Benioff Children's Hospital uses VR to do patient consulting with children who have brain tumors, which improves the patient's experience. That kind of usage gets patients to relax. There, Dr. Kurtis Auguste, chief of the department of surgery at UCSF Benioff Children's Hospital Oakland, gives children and their parents a fly-through of their brains. He can zoom in on their tumors and the nerves and blood vessels surrounding them and explain the procedure he will use on them. This gives the patient much more detail than can be provided by a standard MRI or CT scan.

"This technology is unlike any visual platform I've ever experienced—it is an actual immersive experience," Auguste reports on the UCSF website. "I can use this to plan a surgery and determine the best roadmap to a target in a way I never could before." The system he uses was created by Surgical Theater, Ohio.

It also helps him see things that he would have missed with earlier technology. It's a win-win, and that pre-surgery visualization is bringing us this new kind of "Virtualist" specialist. We will see if the term sticks, but the healthcare practitioners we talked to agree with him that VR and AR are showing up in many places inside hospitals, and they can see a role for doctors and nurses who work only with virtual tools.

These “Virtualists” are the ones setting up 3D sensors and putting HoloLens and other wearable technology into surgery rooms, and the ones who are using VR to relieve pain in burn clinics and using VR and AR to treat phobias, phantom limb pain, eating disorders, and other things. These include surgeons like Dr. Shafi Ahmed, who uses a variety of devices to teach surgery and enable new kinds of surgery to be done.

Dr. Ahmed, who the press has anointed “the world’s most-watched surgeon” due to his broadcasting of surgeries on YouTube and other social media, told us about his work and the work of others who are leading the way to democratize surgery and healthcare by making better education available to surgeons. He wears devices like HoloLens during surgery, which lets him share footage with others, but he says the changes coming are far deeper.

Other places the “Virtualist” will be used are in training, putting new brain research to use, and looking for ways to make the healthcare system more efficient as more Spatial Computing glasses (along with the AI running it all) and robots are brought into hospitals and labs.

Technology coming from companies like MediView, Dr. Ahmed says, will let surgeons actually plan for and perform surgery in whole new ways.

Fixing Humans—Surgery Preparation and Surgery

The crown jewel of many hospitals is its trauma center. Here, patients arrive with various injuries or illnesses, and many are quickly operated on in high-tech surgery rooms. There are dozens of machines and teams of doctors who run around working on saving lives, which is very similar to what you see on TV shows like *Chicago Med*.

What they don’t show you on TV is the preparation that surgeons go through for complex procedures. That used to be done by calling other surgeons who might have dealt with a similar condition. That led to incomplete learning and there being no way to practice a procedure over and over until the surgeon gets it right.

Today, surgeons like Dr. Shafi Ahmed are using VR to view the data coming in from MRIs and other scanning systems, planning out their complex moves, and then practicing them on 3D-printed “virtual” organs with either VR or AR glasses before they get into the surgery room. In this way, they can call in other surgeons and ask for help learning a complex move that they might never have done before, and they can practice it over and over until they get it right.

Then, when they move into the surgery room, new systems show them the scanned data wrapped right onto, and in some cases into, the patient's body. One system from MediView in Cleveland even works like a guidance missile system. It tells the surgeon if they have their cutting tool in the right place and at the right depth with a series of video and audio prompts that are displayed on a HoloLens.

A company in Ohio, Surgical Theater, gives surgeons a very detailed view of, say, a patient's brain, and the surgeon can plan out a path to a tumor, or a group of cells causing epileptic seizures, and then show the plan to the patient and their family.



Figure 8.2: Photo credit: Edward Miller. Dr. Shafi Ahmed using VR to prepare for surgery.

In addition to using VR for surgery preparation and during the surgery itself, Dr. Ahmed has made a name for himself by using VR to train other surgeons. “When I train, what other surgeons need to know very quickly is depth perception,” he says, while extolling its benefits for showing many other budding surgeons how to do a new procedure. Getting a cutting tool into the right place inside the human body often takes great skill, and that learning used to be passed along from surgeon to surgeon. “This takes years and years of experience. What Spatial Computing allows us to do is see the depth.” This changes how the medical system can share techniques, and the benefits go way beyond surgery. As this technology comes down in cost, you’ll see Spatial Computing surgery spread throughout the hospital system. Why? Dr. Ahmed says that the secret sauce is that VR can improve retention of learning by 80 percent.

Dr. Daniel Kraft, who is the founder and chair of exponential medicine at Singularity University, tells us how he uses tools like Osso VR to train other doctors (he has a long list of accomplishments and was the first to implement a text-paging system at Stanford University's hospital). He uses it along with PrecisionOS, Surgical Theater⁴, and FundamentalVR. Surgical Theater lets him plan surgeries. PrecisionOS and FundamentalVR let him practice surgeries. Osso lets him teach surgeries to others. Currently, a small but growing raft of companies, including Boston Scientific, EchoPixel, and Medivis, provide similar capabilities for planning different kinds of surgeries.

Further, there is at least one company that goes way beyond visualization, collaboration, or simulation to actually help the surgeon do the surgery itself: MediView.



Figure 8.3: Photo credit: Cleveland Clinic. Dr. Charles Martin III, Head of Interventional Oncology at Cleveland Clinic, uses a Microsoft HoloLens to overlay AR images on a patient during surgery.

MediView's technology³ is like a missile guidance system for surgeons. It joins data from MRIs and ultrasounds and overlays it in three dimensions on a patient's body to show where a cancerous tumor is. Then, it even guides the surgeon's cutting tools toward the tumor and tells them when they are in the right place to cut it out. This technology will even save the life of the surgeon. MediView's founder told us he lost a partner to cancer caused by ionizing radiation used to locate tumors with old technology and that he knows of lots of surgeons who have gotten cancer from old techniques, which is what drove him to find a better way to visualize tumors on the surgery table.

MediView has already passed through one patient trial of five patients and is about to go through another one.

Additionally, MediView is focusing its AR surgery system for use on soft tissue surgery, which is a much harder thing to accomplish than, let's say, surgery that focuses on a bone, which tends to stay stable and stationary during a procedure. Another company, SentiAR, uses AR to overlay 3D images of a patient's anatomy during interventional procedures that deal with the treatment and analysis of cardiac arrhythmia, so it does deal with soft tissue procedures, but its system is currently utilized for this narrow use only.

Focus on MediView

We recently interviewed Mina Fahim, CEO and President of MediView. MediView has become a pioneering force in utilizing AR technologies within the medical field, achieving critical regulatory milestones and forming strong partnerships with key industry players such as **General Electric (GE)**. The journey of MediView, from its early iterations to its current state, highlights not only technological innovation but also the challenges of introducing novel medical devices into clinical practice.

FDA Clearance and Regulatory Milestones

One of MediView's most significant achievements has been its success in gaining regulatory clearance for its AR-based solutions. The company currently has two Class I registered commercial solutions and a 510(k) cleared Class II solution. These approvals by the US Food and Drug Administration (FDA) underscore the credibility of MediView's technologies. The FDA's clearance of their Class II surgical navigation system as an imaging modality has been a landmark achievement. The clearance covers the application of AR across various imaging modalities, including ultrasound and CT, allowing its use in procedures from head to toe.

The regulatory process for obtaining these approvals was both arduous and complex. MediView played a key role in organizing and executing the FDA's first summit on clinical Mixed Reality in Bethesda, Maryland, which brought together regulators, technicians, clinicians, and reimbursement professionals. This summit helped educate regulatory bodies about the potential of AR in medical imaging and was pivotal in securing MediView's clearance. Importantly, this milestone marked the first time an AR-based imaging modality had been cleared by the FDA, establishing MediView as a leader in clinical AR.

Partnership with GE

MediView's collaboration with GE is another critical component of its growth and credibility. Their FDA clearance has not only validated their technology but has also elevated their reputation, both within and beyond the healthcare industry. GE has been a key partner in the development of MediView's products, and the collaboration has yielded two main products that cater to three of GE's business lines: ultrasound, CT, and fluoroscopy.

The first product was developed through a partnership with GE and is publicly available on both MediView's and GE's websites. It combines AR-based visualization with traditional imaging technologies to assist in image-guided interventions. Their XR 90 product is a Class II 510(k) system that merges ultrasound and 3D holographic anatomy for surgical navigation, providing clinicians an advanced view of internal structures in real time. This collaboration has resulted in a broadening of applications for AR in medical imaging and has solidified MediView's standing within the healthcare technology space.

The success of MediView's partnership with GE has led to additional interest from other companies within the imaging industry, with multiple groups expressing a desire to integrate MediView's AR-based solutions into their own imaging modalities. The FDA clearance has served as a cornerstone in building this credibility, as it provides external validation of the technology's safety and efficacy.

XR Products and Technological Advancements

MediView currently has three key products under its XR umbrella. The first, XR 30, is a heads-up display that can overlay any imaging modality. The second, XR 50, is an ultrasound-guided system designed for needle-based procedures, which is still in development. The third and most technically challenging product, XR 90, is designed for complex surgical navigation procedures, primarily in oncology, where precision in targeting internal organs such as the liver and kidney is critical.

The XR 90 system has been used in over 40 human procedures and more than 200 total procedures, with prominent clients such as Cleveland Clinic, Mayo Clinic, and Cornell adopting the technology. Despite these successes, MediView continues to face significant technical challenges, particularly around trust and adoption. The wow factor of the XR 90 system quickly shifts to the question of whether it can be trusted in real human procedures. Convincing clinicians to adopt the technology and integrate it into their workflows remains an ongoing effort, requiring extensive training and support.

One of the most significant challenges MediView faces is the technical difficulty of integrating AR with soft tissue imaging. While XR 90 has been revolutionary in providing real-time, 3D imaging of anatomical structures, working with soft tissue remains a formidable obstacle. The ability to dynamically visualize soft tissues in real time, as they deform or shift during procedures, represents the holy grail of AR in medical imaging. MediView envisions a future where AR can provide deformable holograms that adjust in real time based on the clinician's manipulation of the tissue, although this is still years away from practical application.

Expanding into Surgical Navigation and Cardiology

Beyond its initial focus on oncology and gastrointestinal procedures, MediView is expanding into new surgical fields, including cardiology. Given the team's experience in cardiac procedures such as ablation and valve repair, the company plans to develop AR-based imaging solutions tailored for structural heart interventions, where real-time visualization is critical. Their Omny product, in particular, has shown potential for use in interventional radiology, vascular procedures, neuro-interventions, and cardiology, signaling the broad applicability of their platform across multiple clinical specialties.

AI Integration and Future Directions

AI is becoming an increasingly integral part of MediView's XR platform. While AI is not yet fully ubiquitous, MediView has already begun using it to streamline image processing and reduce the time it takes to convert scans into holographic models. This process, which MediView refers to as going from "scan to hologram in under a minute," is a crucial area where AI can simplify previously manual workflows, making the technology more accessible and scalable for clinical use.

MediView also recognizes the potential for AI to enhance clinical decision-making. By collecting vast amounts of data during procedures, MediView's system can help build AI-driven models that provide real-time insights and recommendations during surgery. This, combined with AR visualization capabilities, represents the future of surgical navigation, where clinicians will have immediate access to data-driven insights to inform their actions during procedures.

MediView's journey exemplifies the convergence of AR, AI, and medical imaging, with the company making significant strides in developing technologies that enhance the precision and safety of surgical navigation. Their partnerships, particularly with GE, have been instrumental in gaining credibility and advancing their product offerings, while their FDA clearance marks a crucial step toward wider adoption of AR-based solutions in healthcare. However, challenges remain, especially around the technical hurdles of soft tissue imaging and the broader adoption of AR technologies in clinical practice.

As MediView continues to refine its platform and expand into new surgical fields, its commitment to iterative innovation and collaboration will likely position it as a leader in the evolving landscape of medical imaging and AR.

Companies that have a Spatial Computing system that focuses on surgery on hard tissue such as bone are Augmedix (the **xvision spine (XVS)** system), Scopis (Holographic Navigation Platform), and HoloSurgical.

What is common among most of these company surgery systems is the use of the HoloLens 2 AR headset. The HoloLens 2 offers much better functionality in comparison to the first HoloLens, unlocking far greater potential benefits for pre-surgery and in-surgery Spatial Computing systems.

The improved functionality has to do with an increased field of view, a more ergonomic and balanced design, direct and natural hand manipulation of 3D images, improved natural language speech processing, and direct access to the cloud and AI services from Microsoft.

And some have recently started using the Apple Vision Pro during surgery. In a shoulder joint replacement procedure performed by G. Russel Huffman, MD, MPH of Rothman Orthopaedic Institute, at the AdventHealth Surgery Center Innovation Tower in Orlando, Florida, the Apple Vision Pro was used to provide data and visualizations to the surgical team. In another instance, for a shoulder arthroscopy procedure in Jaraguá Hospital in Santa Catarina, Brazil, Dr. Bruno Gobbato and his team used the cameras of the Apple Vision Pro to view the procedure as it was being done. They had initially used the HoloLens 2 headset, but the much higher resolution and better handling of bright lights made the Apple Vision Pro a better headset to use.

The lesson here is that Spatial Computing is already radically changing surgery and will soon do the same for the rest of the hospital—and eventually for the entire system. For both pre-surgery planning and in-surgery procedures, the aim of Spatial Computing is to reduce the operating time, cost, and complications caused by suboptimal surgical execution. Let's look at the use of Spatial Computing another way—how it actually works with patients' brains to open up new treatment methods.

Brain Tricks/Virtual Medicine

Pain management is one piece of the healthcare puzzle in which Spatial Computing, particularly VR, is showing its most promising results. At the University of Washington, researchers working with Tom Furness and burn victims found that playing in a virtual snowfield was better at managing pain than morphine⁵. You can read their studies at <http://www.vrpain.com>.

This and other discoveries, which we call performing “brain tricks,” have deep implications for healthcare and show us a path to a new kind of healthcare that many can experience outside of a hospital, saving time and stress for patients, saving money for the healthcare system, and solving problems better than the drugs of old.

These new anti-pain brain tricks work through multiple methods—the two most understood are distraction and getting the brain to shut down the nerves delivering pain signals to the brain. Putting light into your eyes is a powerful thing, it turns out. “It is a photonic trojan horse,” Dr. Brennan Spiegel says. “If the brain is anxious or upset, it wants to keep track of pain. If the brain is calm and relaxed, it doesn’t have time to feel pain.”

Dr. Spiegel explained to us how it works. Pain, say, from your finger touching a hot stove, travels from nerves in your finger along nerve highways up your arm and into your brain, passing through several “gates” or “handoffs from nerve to nerve.” When those signals arrive at the brain, he says, it is possible to either distract yourself from those signals, or go further and get your brain to actually close the gates between your finger and brain so pain signals never arrive at all. The previously mentioned snowfield app mostly works through distraction.

When you are having fun throwing snowballs, your brain’s attention centers are more focused on the fun than the pain, but Dr. Spiegel is seeing several other pain-blocking effects. He says VR actually affects how serotonin gets used by your brain and nervous system.



Figure 8.4: Photo credit: Cedars-Sinai. Dr. Brennan Spiegel attends to a patient using VR to diminish pain.

People who play Beat Saber, a VR game, note that they go into an enjoyable “flow state” after playing for a while. The effect that people experience, Dr. Spiegel says, is deeper than just distraction, and that VR works just like opioids: they block those nerve handoffs.

“Many people are not in pain even after the headset gets removed,” Dr. Spiegel says. “That suggests there are other mechanisms than mere distraction at work. Engaging people in a flow-engaging activity is the goal.” So, don’t be surprised if your doctor prescribes a VR headset soon and asks you to play some games similar to Beat Saber!

This set of discoveries has a deep impact on not just the opioid crisis hitting the US, but it provides hope that a whole range of mental issues can be fixed simply by getting patients to put on a headset.

“We are all searching for safer analgesics,” Dr. Spiegel says, explaining that in the US, there are lots of research dollars chasing new ways to deal with pain that don’t require using opioids, which are killing close to 40,000 people per year in the US alone. Some companies that have VR apps that address pain management include AppliedVR, Psious, and Firsthand Technology.

Dr. Spiegel and other researchers say that their understanding of these brain tricks is far enough along that they are starting to ask questions about how to use VR with more patients. How do we pay for it? How do you keep the equipment clean? How do we offer a VR consultation service? How often do you need to use VR to treat various conditions? Who owns the equipment and software to make it happen? The field of VR pain management is developing into a standalone discipline. “Dr. Spiegel said, VR works. This tool cuts across all of healthcare.”

An example of another area where VR is being used in healthcare is dementia. With dementia, the “tape recorder” of the mind has stopped working. Dr. Spiegel says of dementia patients, “They have lost their sense of self. That is disorienting and depressing.” Often, dementia patients develop severe depression. Here, a different brain trick is used called reminiscent therapy. He explained that often he deals with patients who are barely coherent and are extremely depressed by their fractured mind. Pulling out a photo of a childhood memory reminds them of something before their memories stopped working, which makes it far easier to work with the patient to get them to further fix their brains. He says this process works better in VR because you can quickly immerse a patient in childhood memories.

Another area for VR’s use in healthcare is schizophrenia. Schizophrenics hear voices in their head. Most of us hear voices, but we understand them to be ourselves. People with this condition hear other voices, akin to living in their own form of VR, Dr. Spiegel explains: “[These voices] aren’t real. Ironically, it takes Virtual Reality to access their Virtual Reality.” Researchers have found that using avatar therapy helps to climb into the mind of the schizophrenic. Patients are shown re-created faces and minds so they can see their demon—and maybe turn that voice into something that helps them out. Dr. Spiegel says that they are seeing statistically significant results that are much better than the drug therapies currently being used.

Another place where the relaxation technique works is during childbirth. Mothers who wear VR see not just a sizeable decrease in pain but perceive that time moves faster, similar to how a Beat Saber player sometimes loses track of time. Tripp VR, a meditation app, is another good example of using the brain to enhance cognitive flow.

Relaxation isn't the only brain trick that comes with VR, though. We are seeing other researchers unlocking brain tricks that come into play with stroke victims, PTSD sufferers, and people with eating disorders.

During a visit to Stanford University's Virtual Human Interaction Lab, we saw how researchers are using brain tricks to help stroke victims and other people who have lost control of one of their limbs. Jeremy Bailenson and others are doing a ton of research on other VR effects, from environmental education to studying whether it can be used to increase empathy for others.

Inspired by Bailenson, Dr. Sook-Lei Liew, who is director of the Neural Plasticity and Neurorehabilitation Laboratory, also uses avatars to help stroke victims. They have worked for years on a study where they give stroke survivors who can't, say, move a limb a virtual avatar controlled via **electroencephalography (EEG)** signals gathered by sensors placed on a patient's skin. Dr. Liew's found that can help their brain and nervous system to rewire, which would eventually help them to move that limb again.

In addition, Bailenson told us about studies that use VR to "mirror" a healthy limb onto their other limb and, often, patients regain some movement and control of the limb that previously didn't move. VR actually can cause new brain pathways to be built, it appears to researchers. This research is exciting for all sorts of other conditions, like PTSD or autism.

"It's very easy to fool the brain," says Dr. Albert "Skip" Rizzo, while explaining his work on using VR exposure therapy to treat PTSD. Dr. Rizzo is the director of Medical Virtual Reality at the **Institute for Creative Technologies (ICT)** at the University of Southern California. He takes patients into a traumatic event, like a simulation of a battle for a veteran who is suffering. He and clinicians can use that visual to trigger them to talk about what happened, and then they can treat it. He reports that VR produces remarkable results.



Figure 8.5: Photo credit: University of Southern California Institute for Creative Technologies (ICT).

Dr. Skip Rizzo served as the project leader for the ICT's Bravemind project, which incorporates VR as a part of exposure therapy to provide relief for post-traumatic stress. The preceding figure is a still from the VR experience.

Dr. Rizzo started the Medical Virtual Reality lab in 1995 after getting excited in the early 1990s about an experience that opened his eyes to the possibility of game-based rehabilitation. He found that a 22-year-old man who had injuries to his frontal lobe made improvements in cognition after playing video games.

VR could also help with autism. In autism, the cells that deal with the ability to correctly discern emotions, for instance, haven't formed properly. Spatial Computing can solve that problem with brute-force methods. Someone with autism might have glasses on in the future that simply say: "Your friend is sad." However, studies are underway to see if some of VR's brain tricks can be used to actually get patients to build new pathways to understand when their friend is sad even without being told via glasses. Dr. Rizzo is also working on reducing anxiety in people who have autism, with a combination of VR experiences and video. The primary scenario that is being used is that of a job interview, where patients can practice going through an interview over and over until they are confident and relaxed.

Another brain trick is feeding it some education or changing the perspective of a patient. This works particularly well on eating disorders. A little education in VR goes further than other approaches, several researchers told us. How can they get patients to eat less salt, which would lower their blood pressure? It turns out that by showing a study group how their body works in VR and what happens as salt begins to destroy their organs, they ate less salt and saw an average seven-point drop in their blood pressure, Dr. Spiegel reported to us. Other studies are seeing an effect if food sizes are virtually made bigger or smaller. If your glasses make those Oreos look bigger, you will eat fewer of them.

Another method of dealing with eating disorders includes getting patients to see their bodies differently. Often, when presented with someone who isn't eating enough, it is because the patient sees themselves as too fat, like what happens with someone who suffers from anorexia nervosa. Researchers are using one of the magic powers of VR: embodiment, which is where you can become something you aren't, either another person or, as Chris Milk, founder of Within, showed us, various animals. In the case of eating disorders, they take patients into the bodies of normal people so they can see that they are sensing their bodies improperly. This remaps their brains so that they no longer feel fat or skinny, and start eating better.

These brain tricks work with a whole range of conditions, Dr. Rizzo reports, from Parkinson's, where VR's brain-relaxing conditions help, to women trying to avoid birth pain. Dr. Melissa Wong, of Cedars-Sinai, uses them when performing difficult procedures on pregnant women. She showed attendees at the Virtual Medical conference how one procedure had women screaming in pain, but with VR on, the patient never winced and was more bothered by the game she was playing being interrupted.

Additionally, researchers like Bailenson and doctors like Dr. Rizzo say they are seeing amazing results while the patient goes through rehabilitation too. VR gets these patients to not only move more, say, when stretching an arm that has been operated on, but do so more often. Many of these exercises are to be done at home, and researchers say they are seeing better results because exercising in VR is more fun, something many VR users are discovering after they buy a VR headset and start playing games that require physical activity. Robert Long, a VR fan, reported on Reddit that he lost about 200 pounds mostly by playing Beat Saber.

Dr. Diane Gromala does chronic pain research at Simon Fraser University and found that adding biofeedback loops improves results even more. Changing the scene you are in or the avatar you are interacting with in relation to, say, your heart rate, lets you reduce pain levels to an even greater extent.

We tell you all these stories (which are only a small fraction of the good news being reported) to show that a new kind of healthcare is coming: one that uses Spatial Computing technology to solve a wide range of health issues. The techniques being discovered here could be delivered to patients as they sit at home, or even as they move around in the real world. How? We see a new kind of virtual coach that will help deliver this new “virtual” healthcare to you. Add that to the “Virtualist” who works inside the healthcare system, and you’ll see healthcare administrators get excited as they see costs come down and fewer patients experiencing negative results due to addiction to opioids. This can only be a good thing.

We’ve seen how VR can be used to fix problems with the brain; now, let’s look at how it could be combined with other applications to give us a “check engine light” to further improve our health and healthcare.

The Virtual Coach and the “Check Engine Light”

When Meta founder and CEO Mark Zuckerberg goes running, he has a real coach. You see this coach running next to him in photos he uploads to Facebook. That coach pushes him and keeps him company on runs, and probably doubles as a security guard. Most of us can’t afford such a coach, which is why so many of us join fitness classes at gyms. Now, however, a new virtual coach has emerged on our watches that cajoles us into doing more steps or performing mental health tasks like remembering to take a few deep breaths once in a while. You see that virtual coach in today’s Apple Watch, for instance, as a circle of lines. It tells you that you haven’t taken enough steps today, or it’s time to take a break from sitting in front of a monitor.

When we’ve run with a real coach, that coach often rode a bike next to us, pushing us on runs, or getting us to do interval training up hills. A virtual coach could easily take over those roles and more.

This new virtual coach will appear in your Spatial Computing glasses. A virtual being might run along with you, pushing you to keep ahead of last week’s run pace. This coach will have access to all your data, or at least the data you want it to have access to. It will have access to your location, gait, and heart data at a minimum, and will know whether it is pushing you hard enough to reach your fitness goals.

It will go a lot further, though, in motivating you to do better. On YouTube, you can see a video of 6,000 people running across Central Park in New York to catch a Pokéémon. That wasn’t a very compelling visual when users of Pokéémon Go stared into their mobile phones, but it shows how computing visuals can be used to motivate people to exercise⁶. With glasses that bring us AR, your healthcare provider will go a lot further, we predict—helping you choose the right foods, smoke fewer cigarettes, exercise both body and mind, and more.

This leads to lower healthcare costs, healthier lives, and more information for the coach to have a chat with your doctor about, if you allow such a thing (and you probably will, given the benefits). The games it will play will make that early *Pokémon Go* game seem quaint soon. Already, we know of people who have lost many pounds playing *Beat Saber* in VR. Imagine games that can improve your eyes, like the one that Calder has built, or exercise your body, or get you to eat better. These are just a few examples of how healthcare is about to get more intimate, personalized, and in many ways “built into” daily life.

Research already shows that if food portions are made bigger in either VR or AR, you will eat less of that item (the research study, funded by the University of Tokyo, used a variety of foods in its tests to show participants different serving sizes). Imagine a virtual coach who assisted you with your food goals this way with you all day long.

Dr. Ahmed explains that this kind of social reengineering is bound to happen and that patients will be persuaded to take the journey into a preventative world. “Your glasses will look at your food and the AR will show you calories,” he told us.

“Getting people to change is the holy grail,” Dr. Spiegel says. “Human beings understand concepts better when presented with more direct evidence, such as that which is visual and made more accessible by being three-dimensional.” Translation: a virtual coach could get you to learn and change faster, which will lead to massive changes ahead for everyone in the healthcare business, and more fun for patients as their virtual coaches entertain them while exercising or eating out.

These kinds of virtual coaches might not look like anything human, either. If you visit a new gym in San Francisco or Boise, Idaho, run by Black Box VR, they have weight-lifting machines that you use while you are participating in a VR experience. If you pull harder on the cable coming out of the machine, the visual in the headset plays a new kind of game with you. As you work out, the machine increases the resistance as you “level up.” We tried it and it uses all the pain distraction techniques that the University of Washington used with burn victims. While exercising, we didn’t notice how hard we were working while playing a variety of games, for instance. Most of the games are things like punching out a monster. Once we were out of the game, though, we were sore from the day’s workout. On return visits, it remembered us and pushed us to go further. The game promises, “Every minute is optimized through AI and adaptive workout prescription.” It uses visuals so that your brain perceives danger—virtually—and gets you to work out harder to achieve your goals. The AI even watches your workout so it knows not to push you past injury levels.



Figure 8.6: Photo credit: Black Box VR. A new kind of gym you use while wearing a VR headset.

This is the age of virtual coaching and preventative healthcare. To enable all this, we need a ton of data to drive the future of personalized, predictive healthcare.

Data, Data, Everywhere

In 2024, as we write this book, there are still many hospitals around the world that haven't digitized their patients' health data yet, and even those that have are struggling to make these **Electronic Medical Records (EMR)** systems useful or the data in them clean and understandable.

Doctors aren't famous for using clear, consistent language, and they often use tons of shorthand and very specialized terms that can contain lots of typos upon entry. We heard tons of stories about problems when we attended Ray Wong's influential Healthcare Transformation Summit for CIOs of healthcare organizations about how medical systems are struggling to not just get the data in these systems into a useful form but get doctors back to focusing on care rather than inputting a ton of data that doesn't lead to better healthcare for patients.

We see Spatial Computing glasses as a potential answer. It is in such supporting technologies that the solution lies, whether to listen to doctors' and nurses' intent and language, support new kinds of computer vision that can look at various things and come up with better data inputs, or connect all the systems in a hospital to talk to your health record. The pressure to improve the efficiency of the system so more people can afford good healthcare will push administrators to continue investing in new technologies like the ones we've laid out.

How many hospitals still have a ton of machines that don't talk to each other in a patient's care room? Most, if not all. All of these problems give skeptics the encouragement to say it's impossible to get to the world we are laying out, and other skeptics are worried the surveillance systems that are under development will lead to dystopian outcomes (your glasses could soon see that you eat too many donuts and could increase your insurance costs accordingly, they worry, and their fears get worse from there). We will take on these fears, and the balance between utility and dystopian control, elsewhere in the book.

Many parts of the healthcare system will be slow to upgrade, like changing a simple drug dispensing machine that might take years to transition to a new always-connected digital model. Connecting machines to a cloud is necessary for the complete overhaul of the medical system, but Douglas Purdy, who used to be a director of product management, gave us another place to look at just how deeply data and machine learning could affect healthcare and all of our lives. He did this by showing us that data is already changing how we can provide healthcare in ways that were hard to dream about only a few years ago.

Purdy built an AI that monitors heart rate variability (which is the signal in between the electrical pulses that pump your heart) and told us that if you train a computer correctly, it can detect patterns in the information between your heartbeats that can tell whether or not you are getting enough sleep, or taking care of your mental health, among other things.

He explained that the heart doesn't actually pump all at once, which is why there's so much untapped data here. What actually happens is an electrical signal starts in one nerve pathway in your heart, causing the muscles to contract, and that signal passes literally cell by cell through your heart.

Humans simply can't watch this, in the same way that humans had a hard time telling whether a horse's feet totally left the ground when running until motion pictures came along. But now, our sensors are getting so good at recording the heart that new patterns are being found in between heartbeats. AI can see those patterns where humans can't.

Purdy is using that AI inside an app called Zendō⁷ to watch your health and tell you if you are meditating correctly, or enough, among other things that are on the way, including sleep, diet, and more. Dr. Spiegel and others validated that heart rate variability is something new and an exciting place to use AI that has only recently become usable for consumers because of the widespread use of watches that have heartbeat sensors built into them, and they said there's exciting new research that is underway now that we have powerful computing tools to analyze patient data like that in real time.

Even Apple now has a Research app that you can download and use to add your own experience—data, essentially—to a variety of medical studies underway. This, along with other data, is doing something new, and could lead to the “check engine light” functionality discussed earlier. Will it lead to your glasses telling you in advance that you will soon have a heart attack? That’s a bit presumptuous for us to say today, but we can see how data will lead to that kind of system.

Dr. Brandon Birckhead put it clearest in our discussions: “A new era of precision VR therapy is here,” he says. “We will be able to tailor things precisely down to the person.”

That kind of hyper-personalized, intimate healthcare is possible thanks to wearable devices like Spatial Computing glasses along with watches and other devices, like a pacemaker that Boston Scientific developed that listens to your heart directly, leading to benefits. Said pacemaker can save your life if you aren’t following doctor’s orders to get enough exercise, or are eating unhealthily, for instance.

Dr. Rizzo told us he is seeing in his studies with PTSD sufferers that many actually share more when they think they are talking to a computer. “They self-disclose more information,” he says. “Contrary to popular belief, people are less concerned about shame [when talking to a computer],” he told us. They worry less about consequences. Even spoken words contain data that can be quite revealing. At SRI International, researchers showed us an AI-driven system where you read a few hundred words from a card, and it identifies whether you are suicidal or not. Stress in your voice can’t easily be detected by other people, but a computer can see patterns most will miss.

Dr. Daniel Kraft explained the role of data best, and it probably comes from his work in the military; he was a flight surgeon but did a bunch of other work, including conducting research on aerospace medicine that was published by NASA. He says all this data will lead to a new kind of situational awareness machine that will watch all parts of the healthcare industry.

Today, he says, you are in the dark with a lot of healthcare. From staffing, to triage, to people pushing food trays around, inefficiency abounds within the hospital infrastructure. Worse, this could be compounded by poor communication, and in a critical environment such as healthcare, this can lead to tragic results in the worst-case scenario. He sees a lot of misaligned incentives, which lead to bad outcomes, but this “situational awareness machine” could exponentially learn as doctors, nurses, administrators, hospital staff, and patients themselves teach it new things and correct old, bad patterns of work.

Dr. Rizzo sees that all this will build a new virtual environment that will help us in every step of our healthcare journey. He says: “The future of this is a companion, buddy, or support agent. Medicine is always about dose/response, and if you can quantify those things better, without compromising patient privacy, we are going to find relationships, or causal things, that go beyond our brains’ ability to read journal articles and learn with traditional techniques.”

Precision and Personalization

Data collection will only increase as time goes on. Our modern-day society is obsessed with precision and personalization. All of this data collection combined with AI and Spatial Computing is bound to produce healthcare that saves many more lives than it does now.

Spatial Computing is transforming not only the healthcare system in labs and hospitals, but is also paving the way for a new era of personalized medicine. This technology can help individuals stay healthier and avoid the healthcare system for longer periods, and once they do enter it, it enhances their care experience. Through AR and VR headsets, patients can now receive treatments and monitoring that were unimaginable just a few years ago, offering entirely new ways of delivering care. This brings us to the end of our journey through healthcare and the changes that are expected to come. In the next chapter, we’ll take a look into a very different world—banking and trading—and how Spatial Computing is also making waves in this sector.

References

1. Exercising in the metaverse with Quest 3—Nerdywork (YouTube): <https://www.youtube.com/watch?v=g8XRpHyHSZw>
2. Sherylle Calder explains EyeGym basics—News24 (YouTube): <https://www.youtube.com/watch?v=4k7cdLNoBcI>
3. MediView XR Technology: <https://www.youtube.com/watch?v=44bo22Kinyk>
4. Precision Virtual Reality™ Montage—Surgical Theater (YouTube): <https://www.youtube.com/watch?v=8DnEadLgWCE>
5. Immersive Virtual Reality pain distraction—World VR Forum (YouTube): <https://www.youtube.com/watch?v=DLTSbhwSQkM>
6. *Student loses 140 pounds in a year playing Pokémon Go*: <https://www.cultofmac.com/news/student-loses-one-forty-pounds-pokemon-go>
7. Zendō: <https://innerflow.es/studio/zendo/>

9

Vision Six: Virtual Trading and Banking

The future uses for Spatial Computing in the financial industry are vast. There is very little Spatial Computing currently being utilized in that field, however. The main reasons for this include the general conservatism of the industry and the number of regulations that would need to be addressed relating to customer interfaces and trading. We believe that the great benefits that Spatial Computing could bring to the financial industry will enable the hearty adoption of its technologies in the future.

Since almost all of the views on the use of Spatial Computing in the financial industry included here are forward-facing, this chapter will be much shorter than the other six chapters we have written that address industry verticals. The functional areas we touch upon here include 3D data visualization, virtual trading, ATM security and facial payment machines, and virtual branch functionality and customer services. Let's start off by discussing 3D data visualization.

3D Data Visualization

One area where a start has been made is in data visualization. A company that has focused on using Virtual and Augmented Reality as mediums to visually present data is a start-up called Virtualitics, based in Pasadena, California. Many of the management team and employees went to Caltech; their understanding of data science is extremely good.

A major focus for the company has been to cater to financial clients who need to have complex data represented in a more visual way in order to simplify finding the possible solutions they are looking for.

However, the software that Virtualitics has built is not specially built to handle certain problems or industry verticals—it is very open-ended.

Another company that had a 3D data visualization product, albeit at a very early stage, was a company called Looker. In 2019, Google announced it was acquiring Looker for \$2.6 billion.

There are skeptics we have spoken with who are very enamored with the uses of Spatial Computing, but they do not see the usefulness of an open-ended 3D data visualization product. Since there are no companies that specialize in 3D data visualization products just for the financial industry, a big opportunity exists in this area. Virtual trading, that is, trading that would happen with a person using a **Virtual Reality (VR)** or **Augmented Reality (AR)** headset, cannot happen without these kinds of specialized 3D data visualizations. Any company, whether it be a bank, a hedge fund, an over-the-counter options and futures trading house, or any other financially oriented company that wishes to do virtual trading, would need to have programs that have 3D data visualization and blockchain capabilities integrated into them.

Currently, many kinds of trading that are not automated are accomplished only after a trader has looked at several sources of information and come to a decision. Most of the time, these traders have to view all the different kinds of information across many screens as simultaneously as possible. Due to this complexity, the rate of human error can be quite high. With the right kinds of 3D data visualization software capability, the numerous screens could be reduced to just one 3D representation, which would enable a trader to make a faster, more informed decision about a particular trade.



Figure 9.1: Photo credit: Bloomberg. A typical hedge fund trader's desk.

Currently, traders tend to use 19-inch monitors, in multiples of 4, 8, or 12. What is interesting is that there are still innovations being developed for traditional computer monitors from companies such as Asus, Samsung, Acer, LG, and Dell.

However, with 3D data visualization software, multiple charts (literally an almost limitless number of charts) could be viewed and manipulated in both Virtual and Augmented Reality. These could be placed side by side, tiled from top to bottom, tiled backward and forward in space, and in combinations of these at different angles in space. Also, AR and VR would be good for providing alerts for specific relevant financial data and important news. This is what the Apple Vision Pro does with ease.



Figure 9.2: Photo credit: Dell. Dell's 43-inch monitor geared toward traders.

In addition to virtual trading, 3D data visualizations would also be useful in portfolio management since the intersection of many different kinds of information is often needed to make a decision there as well.

J.P. Morgan Asset Management, in partnership with Coffee Labs, has launched a new mobile AR experience, providing clients with an interactive way to engage with crucial economic insights that influence markets and investors. This AR feature is incorporated into the acclaimed **Guide to the Markets** and allows users to project economic analyses into their own environments using just a smartphone or tablet via a web browser.

In this immersive experience, users can interact with a holographic representation of Dr. David Kelly, the Chief Global Strategist, surrounded by three-dimensional charts and data tables, enhancing their understanding of market dynamics.

Since much of what could be applied to portfolio management could be used for virtual trading, with even more specialized visualization being used there, we will focus on virtual trading in the next section.

Virtual Trading

There has been much talk among tech and financial insiders about how Spatial Computing, especially Augmented Reality, could be applied to virtual trading, but very few tests so far have been done with actual traders. Both Bloomberg and Reuters experimented early on with initial mock-ups for trading desk software using Google Glass and then, more recently, using a HoloLens headset. The Apple Vision Pro could very well start to solve this issue, as well as future iterations.

In 2016, the consultancy 8ninth (now a part of the Valence Group) was engaged by Citibank to provide a proof-of-concept mock-up for a trader's workstation that utilizes AR and a HoloLens headset. This proof-of-concept became very popular, with many believing that an actual software app had been developed, but it had not been.

Visual aspects of 8ninth's proof-of-concept mock-up include 3D bubbles of individual trading vehicles (these could be stocks, options, and so on) of a given market. The idea is that when a trader clicks on a bubble, they will be able to access its up-to-date trading information. 3D tabs appear to be able to be tiled in space and could provide all kinds of other information, including news. There is also a space for 3D data visualizations.

What is not evident in the mock-up is how a trader would be able to execute a trade using AR software and an AR headset. It also appears to us that there could be consolidation among the kinds of information that an individual 3D object could offer. For example, clicking on a 3D bubble could present both trading information and access to news, as well as other relevant information, such as access to a list of a financial vehicle's competitors that could be clicked on from there.

What does work very nicely is the idea that the size of a 3D bubble, as well as the color of it, could indicate different discrete features, such as the relative market size and industry vertical, respectively. Other novel ways could be integrated into the user interface so that many other types of features could be easily visually evident.



Figure 9.3: Photo credit: 8ninth and the Valence Group. This is a proof-of-concept workstation mock-up of what a trader could potentially access using AR and a HoloLens headset, developed for Citibank.

Besides wearing an AR headset, another possibility was to wear smart contact lenses made by Mojo Vision, called Mojo Lens. Mojo Vision, a Saratoga, California start-up with dozens of PhDs and engineers, came out of stealth mode in January 2020. Since having a viable version of the Mojo Lens would take several more years and quite a bit more capital, Mojo Vision pivoted its focus to the production of micro-LEDs in 2023. Using an AR headset and specialized software, all kinds of trading support could be accomplished, which includes equities, options, futures, **real estate investment trust (REITs)**, commodities, and any other tradable financial instrument. Where the most benefit would be gotten, though, are those instruments that require the analysis of many different streams of information in a fast manner before making a decision to trade. These include options and futures, both at times when decisions have to be made as to whether or not to purchase the instrument and also as to whether or not to “roll forward,” that is, to extend the expiration or maturity of an option or futures contract by closing the initial contract and opening a new longer-term contract for the same underlying asset at the then-current market price.



Figure 9.4: Photo credit: Mojo Vision. Mojo Vision's Mojo Lens feature roadmap for the 2020s.

The amount of information needed for these decisions would normally necessitate the use of many monitors. The stress of having to consult these monitors and then have to make very quick decisions as to what to do has worn down many traders. Before reaching trader burnout, traders can make very costly errors. Using an AR headset with software that supplies 3D data visualizations with plotted relevant points based on criteria input previously could cut down on burnout and errors. News and information could be pulled up alongside the 3D charts, with elements there added as criteria for the charts themselves. Actual trade order execution could also be achieved.

Some delay in using AR headsets and specialized software for information and decision-making, as well as trade execution, will come from financial regulatory authorities. Getting their “OK” will take time because of the potential risks involved. If information from Reuters, Bloomberg, or any other company that provides financial quotes and news is not updated in exactly the same timely manner in the new software as it is in existing systems, then the new system will not be reliable. And if trade execution lags by even a millisecond compared to existing systems, there would be a major issue¹. The new software would have to match with existing systems when running tests; it would have to match at all times in every instance.

Having said this, there is no technical reason why systems using AR headsets would not be able to match existing systems’ reliability. We believe that the future of trading is one that relies on Spatial Computing.

Management from the New York Stock Exchange, Nasdaq, and the hedge fund Bridgewater Associates expressed interest in our book but declined to comment. We know of at least one very large global financial services firm that has a great interest in Spatial Computing and has started developing technology for it, but we have been asked to keep it confidential.

Different aspects of Spatial Computing, including Computer Vision and associated AI, make the next section we are focusing on in this chapter possible.

ATM Security and Facial Payment Machines

It is now commonplace to unlock smartphones and computers using fingerprint and/or facial recognition. In some countries outside of the US, such as China, Spain, and South Korea, facial recognition has already been implemented for ATM security.

When we spoke with Brett King, the author and founder of the banking service company Moven, he became especially animated when talking to us about the developments happening in China. In China, facial recognition is even being used to pay for vending machine items.

Chinese technology expert Matthew Brennan has reported that he found a TikTok video that shows that non-Chinese people are blocked from using that facial recognition payment system, presumably because they don't have IDs issued by mainland China.



Figure 9.5: Photo credit: TikTok. A still from a video uploaded to TikTok that shows a woman paying for a vending machine item using facial recognition.

In 2015, China Merchants Bank launched an ATM that uses facial recognition, eliminating the need for the use of cards. It was the first time that this had been made available for public use. The company claimed that the system would be able to distinguish between twins since several parts of a face were used to create the profile used for facial recognition comparison.

Wearing glasses or different makeup was tested and shown to not make a difference.



Figure 9.6: Photo credit: China Merchants Bank. An advertisement for China Merchants Bank that tells users they can now use their face to withdraw money.

Several banks in mainland China, including the Agricultural Bank of China, China Construction Bank, and China Merchants Bank, now have ATMs that use facial recognition as part of their security process.

Taiwanese bank E.Sun partnered with NEC to provide ATMs that use facial recognition; the ATMs were made available to the public in February 2019. In order to set up facial recognition cash withdrawals, customers must first insert their cards and then have a photo of their face taken with a camera installed within the ATM. They are then sent a one-time password to their mobile phone, which must be entered into the ATM within a minute, thus registering their facial image.

In Spain, in February 2019, CaixaBank started offering ATMs that use facial recognition. Trials conducted among customers before indicated that 70 percent would use it as an alternative to entering their personal identification code on a keyboard. Registration for the facial recognition system is conducted in a bank branch and takes a few minutes using tablets.

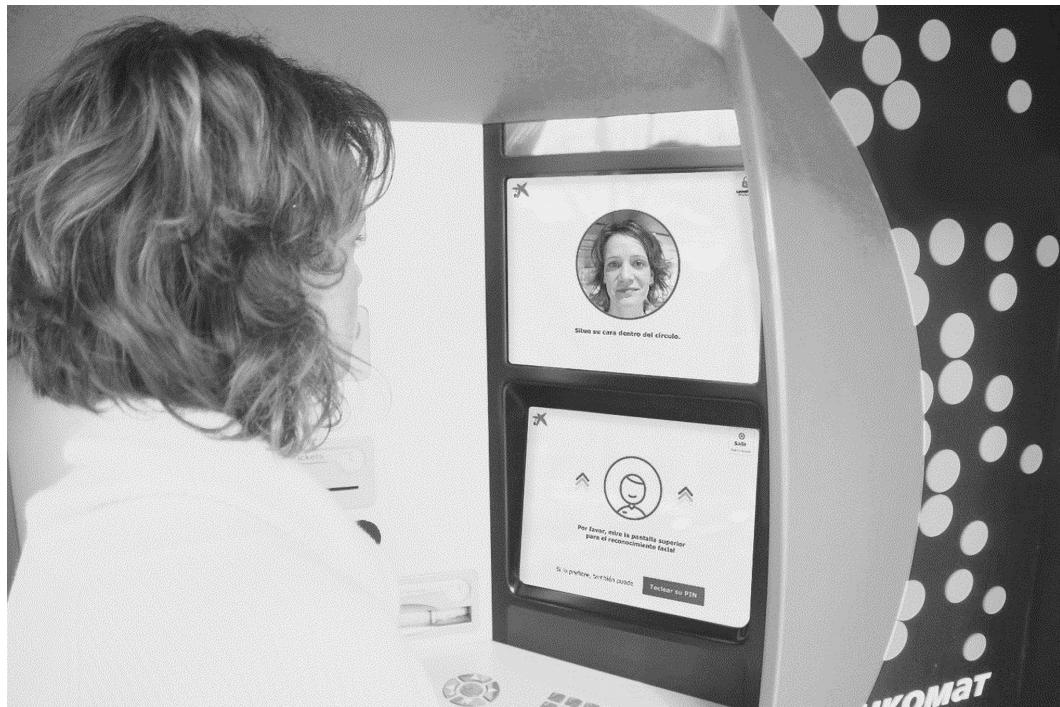


Figure 9.7: Photo credit: CaixaBank. A CaixaBank ATM that uses facial recognition instead of a card to access its system.

In 2023, South Korea's Shinhan Bank rolled out the country's first biometric facial recognition payment service Shinhan Face Pay to 900 ATMs. South Korean start-up Alchera aided Shinhan Bank in these endeavors.

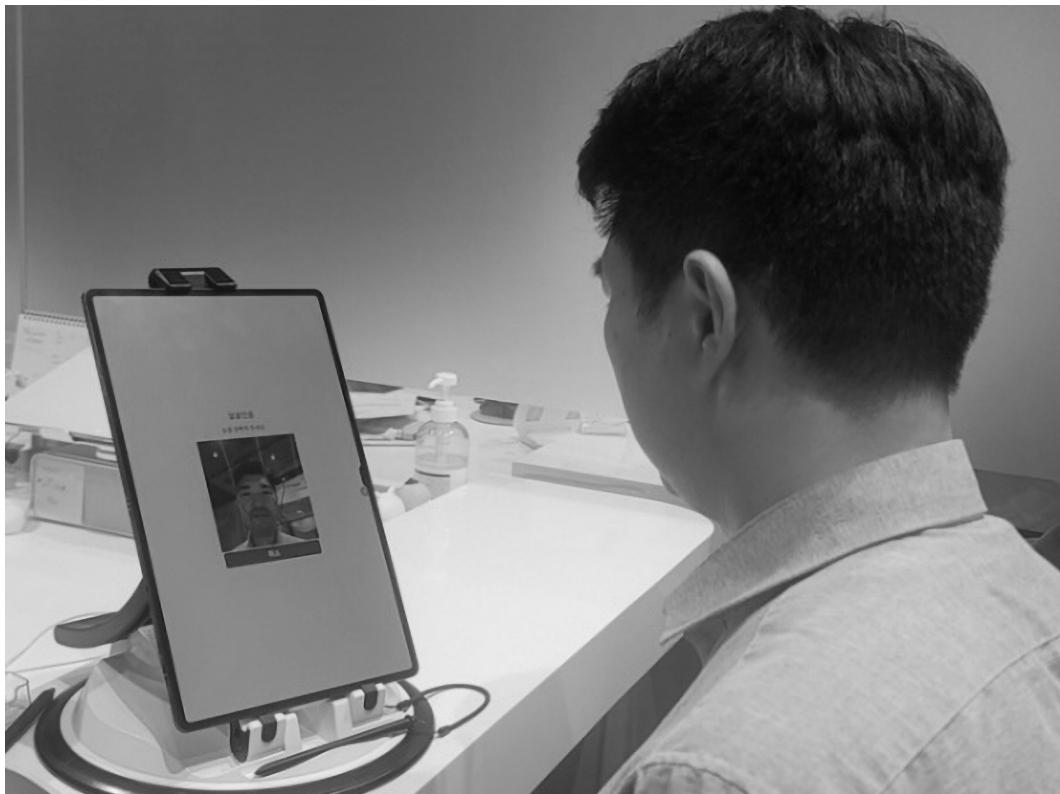


Figure 9.8: Photo credit: Shinhan Bank. A man demonstrates how Shinhan Face Pay works.

Outside of banking and ATMs, facial recognition could be used to pay for products. Again, in China and in Korea, this has already been achieved with “facial payment machines.”

Mobile payment in China is already the most advanced in the world, with customers making a purchase simply by posing in front of **point-of-sale (POS)** camera-equipped machines, after linking an image of their face to a digital payment system or bank account².

With devices already in over 300 cities, Alipay, the financial arm of e-commerce giant Alibaba, has been leading the facial payment machine business in China. Its facial recognition system comprises a screen that is as about as large as a mid-sized iPad, its AI software and Computer Vision components, and mobile payment software called Dragonfly. Alipay has earmarked \$420 million over the next three years to roll out its technology.



Figure 9.9: Photo credit: Alipay. Pictured here are two Dragonfly facial recognition systems, which are currently actively used in China to make mobile payments.

Additionally, in August 2019, Tencent, which runs the WeChat app with 600 million users, announced its new facial payment system called WeChat Pay. Besides Alipay and Tencent, a growing number of start-ups in China are focusing on the facial payment machine business.

In Korea, Shinhan Bank is behind Shinhan Card's POS camera-equipped machines.



Figure 9.10: Photo credit: Shinhan Bank. A customer at an affiliate store in Seoul utilizes Shinhan Card's facial recognition service.

Facial recognition technologies can also be used to facilitate virtual banking, which is the topic of our next section.

Virtual Branch Functionality and Customer Service

Second Life, the virtual world, shut down all the virtual banks that had been built and operated there in 2008. Now, almost none of those banks were real financial institutions but rather individuals who had deigned to create a bank out of thin air and put it in operation among the denizens who populated Second Life. However, Second Life virtual banks actually had deposits made with real money, and ultimately, they were shut down due to a lack of regulation.



Figure 9.11: Photo credit: Second Life. One of the virtual banks that operated in Second Life before they were all shut down.

Nowadays, talk of a virtual bank refers to a real financial institution that operates without any physical bank branches—basically, an all-digital bank. A good example of this is Chime.

CaixaBank, in 2017, became the first bank in Spain to incorporate the iPhone X's Face ID technology so that customers could access their accounts solely using facial recognition on their mobile devices. Now, several banks around the globe offer this capability.

With many of the major banks around the world closing down physical bank branches due to a lack of customer foot traffic, there has been some movement around the idea that all bank branch functions could be digitized, such as those associated with opening up a first bank account and mortgage paperwork.

Shaun Moore, former cofounder and CEO of Trueface, a company that had worked with top global banks on facial recognition uses that was acquired in 2021, commented “*We are seeing the financial services sector test face recognition as a part of multi-factor authentication for ATM withdrawals, mobile banking enrollment, and mobile account access and transactions. By implementing face recognition as the key step in multi-factor authentication, banks are able to mitigate their exposure to risk and fraud, saving themselves millions of dollars in the process.*”

Along with bank services being available post-facial recognition checks are those available using voice technology, with **United Services Automobile Association (USAA)** being the first to offer limited voice banking. Capital One offers banking by voice via Amazon’s Alexa, on devices like the Amazon Echo and Amazon Dot. Capital One customers can check balances, review transactions, make payments, and more using simple voice commands. Bank of America’s “Erica” works similarly, but off of its own app, versus using Alexa-reliant apps.

What we envision is virtual banking that combines the use of facial recognition and voice commands and is accessed via an AR headset. This would also allow both virtual AI assistants and real bank customer service reps to be tapped by the customer on demand when needed.

Future Adoption

Even though the use of Spatial Computing technologies for virtual trading and banking is at such a nascent stage, we are confident that given the great benefits that Spatial Computing could bring, these technologies will be adopted on a large scale in the future.

Spatial Computing has the potential to make trading and banking functions take less time and make more money—particular outcomes that are shared with other industry verticals, including manufacturing, retail, and healthcare.

In our next chapter, we turn to how Spatial Computing could be used for education and training across several industry verticals.

References

1. *Rigged Wall Street, an interview with Michael Lewis:* <https://www.npr.org/transcripts/297686724>
2. *Iris scan:* <https://arstechnica.com/gadgets/2022/10/apples-ar-vr-headset-will-scan-your-iris-when-you-put-it-on/>

10

Vision Seven: Real-Time Learning

Bill Hill told us how he escaped poverty. He grew up in Scotland in dire circumstances and later went on to invent, at Microsoft, the font-smoothing technique that is used on all computers and phones today to make text easier to read. When he was growing up, instead of playing on the streets like other kids, he went to the library. He told us that he read a new book every day and that helped him be “educated” about how the world works.

In the past few decades, the world has changed. Now, everyone has the world’s libraries at their fingertips. You can talk to a Google Home device and ask it “What is the latest in breast cancer research?”.

Google Home has an answer. Yet most people aren’t “educated.” The truth is our kids don’t know how to be curious. Worse still, there are massive new distractions to absorb their attention. YouTube has more videos being uploaded every minute than most humans can watch in 100 years. Fortnite, Roblox, Minecraft, and many other games captivate children’s imaginations far more easily than books can. This is hurting their ability to be curious while removing any chance of being bored, which is when creativity really comes out.

In this chapter, we’ll talk about how Virtual and Augmented Reality is already being used in schools and corporate training. These technologies can provide the same kind of learning (in fact, better learning), many people tell us.

Experiential Digital Learning

Educators report to us that Spatial Computing technologies are already radically changing how our kids are learning, and they see even more change coming. **Virtual Reality (VR)** can show kids things that only rich kids in the best schools could see before. When we've met educated people previously, part of what makes them educated is that they are well traveled, well cultured, and often know about diverse sets of music, theater, and dance, along with a deeper specialization, whether retailing, farming, engineering, law, or medicine. Getting more kids to have the same opportunities is something that VR and **Augmented Reality (AR)** are showing is not only possible but probable. In VR, you can remotely attend an amazing university class and "feel" like you are there. But remote learning is only a small piece of why education is about to radically change. Once learners have Spatial Computing devices on their faces, they will find that learning itself is radically different from what most of us experience in classrooms. Experiential learning will not only reach kids who haven't had any boredom in their lives. To fill their free time, kids are now constantly on their smartphones or tablets, and many are spending hours a day in VR already.

Even if kids could find the books that would educate them, learning is slow, isn't personalized, and isn't often the best way to learn. Research done by James M. Clark and Allan Paivio shows that learning done in 3D in VR or AR is about twice as likely to be remembered as learning done using books or old-style videos.

Take a John Steinbeck novel that talks about Monterey, California; this could be made much more "real" to today's kids by taking them there. The problem is that getting that kind of education is only open to the wealthiest. When we did a 9,000-mile road trip in the summer of 2018 to take our kids to learn about America (and to do the same ourselves), that trip cost about \$12,000 in hotels, electricity for our Tesla, and a month of our time. That simply is not an avenue open to most people.

But here comes Spatial Computing, with a new way to learn, which we call "experiential digital learning." Want to understand how something works? Or visit a historical site? You can do that in your new head-mounted displays, or even while looking at AR on a smartphone screen. Already we are seeing that kind of learning happening in classrooms around the world and, yet, a bigger shift is coming, one that will let many more kids excel and get the kind of education they will need to fix the huge problems that are appearing, such as climate change or social problems due to opioid addiction or terrorist attacks, not to mention new aspirational goals like doing more space exploration.



Figure 10.1: Third-grade students at Forest Hill Elementary School in Campbell, California used VR in their classroom. They were using Google's Tilt Brush in Oculus Quest headsets to discover new ways to make art (photo credit: Robert Scoble)

Jesse Schell, CEO of Schell Games and principal investigator of *HistoryMaker* VR, remarked in a statement, “The most recently published National Assessment of Educational Progress report (NAEP, 2014) cites only 18 percent and 23 percent of eighth-grade students performed at or above the proficient level in US history and civics, respectively. What’s more, students commonly report they find history to be boring and they lack motivation to understand or remember the content.” Jesse went on, “The generous grant from the IES coupled with the talented *HistoryMaker* VR development team could play a significant role in helping those numbers climb in the near future by offering students an entertaining way to learn about the past.”

Why is VR so important to education? Well, for a few reasons. One, it lets you experience learning as if you have experienced something in the real world. Learning archeology? It’s better to walk around Rome or Jerusalem with an archeologist. VR and AR can do that. Learning chemistry? What if you could visit a virtual lab and experience chemical reactions firsthand? VR and AR can do that. Learning physics? What if you could visit the particle accelerator at CERN and hear a talk with one of the physicists there? VR and AR can do that¹. Using AR and VR, kids could learn about basic finance and how to use money.

No Teachers, Books, or Movies? New Approaches in the Offing

It isn't just the classroom that's about to radically change, though; learning at work or even in your kitchen is about to change due to both VR and AR. The same set of questions apply. What if you could learn to do surgery without having a human patient? VR and AR can do that. What if you want to learn to fix a tractor? AR can do that and do it while you actually have your hands on the tractor. In fact, companies like Caterpillar are doing exactly that. What if you want to learn how to cook a meal in your kitchen? AR can do that. What if you want to learn how to do a job on a factory line? Both AR and VR can do that.

A perfect storm is coming, not just because of VR, AR, and AI/machine learning, but some of that is due to new kinds of schools bringing new approaches, like 42 Silicon Valley, which teaches people to do new jobs. We predict that all four will collide in the 2020s to bring massive changes to education and corporate training and will even offer an answer to the massive numbers of people who will need to be retrained to do new things as automation causes many jobs to go away. Truck driving, for instance, is America's number one job, with 1.3 million drivers. Many of those jobs, if not all, will go away in the next one to three decades. And that was before the job shifts that were completely unanticipated due to the COVID-19 pandemic. Huge unemployment currently exists and Spatial Computing, along with new educational approaches, offers an answer.

42 Silicon Valley, a peer-assisted and immersive school located in a big building near Tesla's main factory in Silicon Valley, is an example of a new educational approach. The school is unique in that it has no teachers and charges no fees. It reminds us of how Uber disrupted taxis while owning no cars or how Airbnb disrupted hotels while owning no properties. It was donated by a French billionaire and is packed with Macintoshes and tons of students, who agree to spend more than eight hours a day there.

The school took a chef who had never coded, Danny Saetern, and got him to the level that he started a VR game studio and, today, he is working as a software engineer at Schlumberger. We visited the school and saw students grouped together, learning from a curriculum and on computers the school provided, but learning mostly from each other. Students, when they hit upon something they didn't understand, asked the other students for help figuring it out. By doing this and immersing themselves on-site for about 10 hours a day, the students learned very fast. Danny, the former chef, told us he learned to program in about a year.

Another example of shifts in the way schools operate can be found at Lambda School, which is a school teaching computer science. Here, the change is in how students pay. Since this school is more traditional, with a teacher walking students through their learning, they pay a portion of the wages that they get when they get a job².

Yet another shift is coming from the start-up sector with entrepreneurs like Nick Smoot in Idaho and investor/entrepreneur John Borthwick in New York who are building “Innovation Collectives,” which have daily classes that teach students a variety of topics (often for free) to build stronger companies around the ecosystems they are working to prop up. At both, we saw classes teaching VR and AR, among many other topics.

Phil Komarny, former VP of Innovation at Salesforce, told us he is seeing these educational shifts even inside traditional institutions like Harvard. He says that there will be new efforts to help upskill a disrupted workforce and is seeing lots of new technology to do it, such as blockchain to track a student’s progress, and sees great advancements in the ability to co-learn, like 42 Silicon Valley does, thanks to AR and VR. Blockchain and new collaborative AR and VR systems, he explained, will be used to build a new learner network that will bring many of the same advantages that, say, attending a school like Stanford or Carnegie Mellon does today. A huge part of the value of going to those schools is the network you build with other students, many of whom will go on to build important companies and technologies.

It is this combination of both technology and new approaches that have us most excited about the future of learning, training, and upskilling. Current workers, students, and future generations will need this because many jobs will be automated and because of major new challenges due to climate change and globalization.

Many of us learn best by doing, or, at least, visualizing what we are learning. How many of us, when reading a math equation, visualized it inside our minds so that we could better understand what we were learning? Watch children learn math, for instance, and they will often count on their fingers. Why? Fingers are physical and can be tracked.

Brainier Learning

Spatial Computing is the most important learning technology humans have ever invented. Why? It goes back to how the human brain learns. We learn a lot faster when we are engaged and when we can visualize. Spatial Computing brings both of those ingredients in big doses. Today, learning something out of a book, say, a chemistry equation, might take an hour. Using Spatial Computing, it might take a minute to grasp the same knowledge—not to mention we will probably learn the concepts better.

In VR, you can be embodied in someone or something else. You can experience life through the eyes of another, which, in studies done at Stanford and other schools, builds not only empathy for those others but as deep a memory as if you experienced the learning in the real world. You can also slow down time and take a different view: walk around a chemical molecule, for instance, or look at the results of a math equation in a way that would be impossible with any other method.

That shift is also joined by another one: learning can be delivered in real time on top of things we need to learn to manipulate. Let's say you need to learn how to take an oil pump out of a tractor. In the "old" world, you would have needed to either read an instruction manual and go step by step or sit in a classroom where you would have memorized the steps. With AR, the steps can be shown to you right on top of the real tractor. In some instances, it can even watch you work and warn you that you are about to do something wrong. We've seen manufacturing workstations in place in China that do exactly that, but over the next decade, that kind of learning will come to a consumer's face. Imagine an IKEA-bought piece of furniture that actually shows you how to assemble it, and watches you work to make sure you did it right.

Tractors aren't the only place we are seeing this new kind of learning. Students are using AR to learn everything from how to do math equations to chemistry to biology on virtual cadavers to learning to drive a forklift to playing football better. Football? Yes, at Stanford University, Jeremy Bailenson's students (who included Derek Belch, founder and CEO of Strivr, a company that trains many people in both sports and corporate jobs) took a 360-degree camera onto the practice field and recorded a bunch of plays while the camera was next to the quarterback. The quarterback then watched those plays at night to prepare for the next day. They found his stats got better. So, they tried it on more than 700 other players and learned that people who initially trained in VR versus using a 2D tablet had a 20 percent decrease in reaction time and were correct more often, too.

Why does this work? Strivr's Belch says that it's partly because it's the same view a player will have on the field, but also partly because when players are on the field, their brains pay attention to things that will do them harm³. When a 250-pound defender is running toward you, it's hard to pay attention to other patterns that are taking shape on the field. However, partly because you can play the 360-degree replay over and over, you can learn to see patterns taking shape and train your brain to see them. You actually learn better. Belch agrees and says, "The effect on your brain says this [technology] is legit. Having the headset on is more effective than 2D learning. This stuff works when people buy in and do it right. There is real research behind this."

The learning goes deeper, though, because the system and observers can see where a player, or, really any user is looking. That way, others who are watching along with the one wearing the headset can see that they are looking at the wrong thing or see why a mistake is being made and fix it right there. These effects simply aren't possible with 2D video.

What comes next with that data is a really interesting play. Strivr is recording data from the controllers that people being trained are holding, along with the sensors on the headsets, which can show all sorts of things, potentially including heartbeats and where the user is looking. That data can be used to train new AI systems that can potentially let managers know who won't work out or who is better suited to a warehouse job than one working with people. Already, Strivr has announced a system that will help decide whether an employee should be promoted, and we can see quite a few other HR tasks that could be done in the future, building on data that comes from tens of thousands of people running through its HR training. This technique has moved to other sports. The New York Mets are using it to prepare for tomorrow's games. If you search Google, you will find lots of examples from soccer to basketball⁴.

Others we talked to agree that a new kind of learning has arrived and will radically change the learner's experience and even bring more people into lifelong learning and retraining that other methods simply can't dream of. Classrooms are already seeing teachers use Merge EDU's Cubes to teach everything from astronomy to frog dissection.

Why does it all work? The brain experiences presence. In other words, it treats VR as real as the real world, and that strengthens neural pathways. In the same way that you can remember, in detail, an event like a car crash or an earthquake, so too can you remember the training you received in VR, while studies show that the same training delivered in books, videos, or slide decks can't be remembered nearly as well.

This has huge implications for training and the simulations that come with training. So, why aren't we doing more of this? It mostly comes down to cost. Overall, VR hardware costs when HTC Vive or Oculus Rift first came on the market ran about \$3,000 and needed a room where sensors could be permanently set up. In 2019, the cost came down to \$400 with the Oculus Quest, and we saw a great increase in the number of companies that started using VR for enterprise uses. The Meta Quest 3 now costs \$500. As costs generally continue to come down (though the current Apple Vision Pro is \$3,500) and headsets get more capable, watch for exponential increases in the number of places we'll see both VR and AR used. The one room that might be more changed than any other by this technology is the classroom.

Spotlight on Cornerstone

In September 2024, the authors of this book had an interview with Kyle Jackson, Global General Manager of XR and Spatial/AI at Cornerstone. Jackson previously founded Talespin (AI and Spatial Computing), a skills mobility and corporate experience platform company, which was acquired by Cornerstone in March 2024.

In 2019, Talespin embarked on a groundbreaking journey, focusing on conversational-based learning with the belief that **large language models (LLMs)** and Spatial Computing would revolutionize industries. Talespin envisioned the integration of AI with Spatial Computing as a major unlock, transforming how businesses interact with advanced technologies. Early on, they developed Copilot Designer, a platform that blended these technologies well before the “Copilot” concept became widely known, and they continuously refined its capabilities.

The company's primary goal was to simplify content creation by merging Spatial Computing and AI, removing the need for complex development processes. This empowered enterprises to create simulations and training tools that combined the power of AI and spatial awareness without requiring extensive technical expertise. Its vision resulted in hundreds of simulations across multiple industries, showcasing how AI and Spatial Computing could support no-code and prompt-based solutions for businesses. While there were challenges—especially in the XR space, where companies often wanted to focus on specific, niche use cases—the company remained committed to creating a scalable platform that integrated AI and Spatial Computing to support a broader range of applications.

A major breakthrough came when the company achieved a large-scale deployment at Accenture, where 60,000 headsets were used to roll out simulations. This project, which heavily relied on Spatial Computing to create immersive environments powered by AI, demonstrated the vast potential of their platform. The combination of conversational AI, spatial awareness, and wearable technology allowed the company to operate on a massive scale. It was proof that their vision of combining AI with Spatial Computing could work seamlessly within corporate infrastructures, setting them apart as a leader in the field. This success caught the attention of Cornerstone, a previous investor, leading to the company's acquisition.

The acquisition by Cornerstone in 2024 marked a new phase of growth for the company. What was once a start-up with a small sales team now had access to over 500 salespeople and 140 million users. With Cornerstone's support, the company is scaling rapidly, expanding the application of its AI and Spatial Computing solutions. However, this also required educating the market on how to deploy and scale these tools effectively.

Moving from serving innovation buyers to infrastructure buyers came with new challenges, but it opened up more opportunities for widespread adoption.

The global surge in AI development, particularly Generative AI, also played to the Cornerstone's advantage. Since 2019, its strategic focus on blending AI with Spatial Computing put it in a unique position to benefit from growing interest in AI across industries. As companies began reallocating budgets toward AI, the company's solutions, which integrated both conversational AI and spatial awareness, became increasingly relevant.

One of the most exciting developments on the company's horizon was a major software release, the result of two years of development. This update focused on solidifying their Generative AI systems, which, when combined with Spatial Computing, allowed for rapid creation of simulations—some in as little as 15 minutes. With Spatial Computing enhancing the immersive quality of these simulations, businesses could now generate high-quality, contextually aware training tools more efficiently. This merging of AI and Spatial Computing capabilities represented a major leap in both fields, making simulation building faster, easier, and more scalable.

As their platform matured, the company worked with a variety of notable clients globally spanning financial services, healthcare, retail, insurance, and many other industries. These clients utilized the platform's ability to integrate AI with Spatial Computing for a range of applications. For instance, one large retailer in Australia leveraged the platform for leadership training, while a major European pharmaceutical company used it to build healthcare simulations focused on doctor-patient and staff interactions, with AI-powered avatars helping to improve communication skills in clinical settings. A global financial services firm applied AI-driven simulations for financial literacy and regulatory compliance, using Spatial Computing to create interactive training environments for employees dealing with complex financial products. Another US insurance company utilized the platform for sales training, relying on AI simulations to enhance consultative selling techniques, where salespeople interact with AI avatars representing potential clients, allowing them to practice and refine their skills in a safe and controlled environment.

Even a US-based mortgage company found unique uses for the platform, building leadership and sales simulations that integrated AI with Spatial Computing to train employees in effective communication and empathy—two crucial aspects in customer interactions. These use cases highlight the versatility of the company's platform, demonstrating how AI and Spatial Computing can be combined to solve diverse business challenges across industries.

The company's work revealed a critical observation: the growing gap between the skills required by modern businesses and the existing capabilities of the workforce. As AI and Spatial Computing continue to play a larger role in business operations, this skills gap is becoming even more apparent. AI agents, powered by LLMs and Spatial Computing, are poised to fill this gap, acting as virtual assistants that could interact with employees, provide guidance within their environments, and even offer career coaching. These AI-driven, spatially aware agents have the potential to revolutionize the workforce, helping businesses overcome labor shortages and skill deficits by providing real-time, context-based assistance. AI agents are basically autonomous AIs that, when presented with a goal or task by the person using AI, go about the job until all is answered or accomplished. For example, when a person calls their health insurer, they could encounter as the first step an AI bot that then hands off the task to other AIs until the person who called gets their answer. Another example would be a person who is asking an AI if a particular car should be bought, given the location, price, season, and other details, followed up by a question as to the five best places to buy that car if the answer is affirmative.

Looking to the future, the company saw the integration of real-world understanding as the next significant step. By combining Spatial Computing with AI systems that could interpret physical environments through cameras and sensors, they envision creating even more immersive and contextually aware simulations. These simulations would allow AI agents to interact with the physical world, guiding employees through tasks in real time. For instance, a retail employee wearing a headset could receive guidance from a virtual expert, helping them with tasks directly within their environment, thanks to the combined power of Spatial Computing and AI.

Cornerstone's long-term vision remains focused on creating AI-driven tools that integrate seamlessly with Spatial Computing to offer scalable, accessible solutions across industries. By continuing to develop AI that can interact with and understand the physical world through Spatial Computing, they aim to help businesses bridge the skills gap and prepare for future challenges. As AI and Spatial Computing technologies evolve, the company is committed to staying at the forefront of innovation, believing that the synergy between AI and Spatial Computing holds one of the more important keys to transforming the workforce and solving complex business problems.

The journey they've begun is far from over, but the integration of AI with Spatial Computing has already proven to be a powerful force for change, one that will only grow more essential as industries continue to evolve. Next, we tackle what the future of education will be.

The Future of Education

How do you learn? Some can learn well from a book. The two coauthors find that they learn far better by actually seeing how something works, or, better yet, building it and teaching it to others. That is why we chose the line of work that we are pursuing, choosing to talk to people who are making the world work and seeing how they do it.

VR and AR bring us a very powerful new education technology that can do just that: give you the ability to see, build, and learn about something in a way that closely maps how we learn. Even better, because it maps to what we already know about the real world, like how to grasp a cup, this learning will be open to far more people, even those who struggle to read.

We observed just this as we watched a four-year-old girl use VR. She learned how to use the controllers, how to open an app, and how to copy a file, all within a few seconds of putting on the headset for the first time. What did she see inside the headset that was nearly as big as she was? She could pick up a cup just like she could in the real world. She could navigate around from space to space very much like she does in the real world. And the visuals she was seeing were being perceived by her brain the same as she sees the real world; the whole process was intuitive.

In VR, all you see is this virtual world but your brain perceives it as being just as real, or at least much more so than, say, a book or 2D video, as the real world. You can “immerse” yourself in what you are learning, for instance, by taking a virtual tour of archeology, or getting a tour of the solar system as if you were an astronaut on a spaceship flying past.

If you study how people learn, like Lisa Petrides—founder and CEO of the **Institute for the Study of Knowledge Management in Education (ISKME)**, an education nonprofit—does, you will be excited and scared by the changes coming due to AR and VR to not just the classroom but for lifelong learning. Some of these changes might mean virtual classrooms move from physical spaces to being wherever a student is, whether in a living room or a coffee shop. Already, we have experienced new classrooms (both as a teacher and as a student) in virtual spaces like Nanome and Decentraland.

Lisa is scared because she sees yet another round of kids left behind on another round of technology due to cost and other factors. She told us of walking through a township in South Africa and seeing electronic smart boards rendered useless because, in some neighborhoods, power is up only a few hours a day.

Even in richer countries, she told us, lots of the educators who come to her yearly conferences are seeing huge problems in student preparedness. Many students, she told us, are born into harsh circumstances and are coming to school unprepared to learn. Think about the student who is malnourished, growing up in a homeless shelter, or has opiate-addicted parents, she said, listing some of the huge problems that her institute faces every day.

She sees that AR and VR, even while being huge accelerants in the speed of learning for many kids, could leave lots of kids behind and, worse, tear open new societal wounds that are already straining. Some kids who now are in special needs programs might be simultaneously helped and left out by such technologies as other kids race ahead with their learning due to new Spatial Computing devices on their faces or in the classroom.

Even where new technology could lead to some answers, educators will be asked to figure out how to deal with new privacy fears as the new Oculus Quest has four cameras on the front of it, mostly to see the real world to do inside-out VR tracking, and so it can see where the controllers are in space and even more sensors and data privacy issues are to come.

While these are serious issues that we feel must be considered, we believe that the potential benefits of VR and AR will lift many kids into being educated, who would have otherwise fallen through the cracks. Many more stories like Bill Hill's will be heard after this technology arrives in more classrooms.

Fears aside, she says, she is excited by the tons of great news and new approaches for educators thanks to these new technologies, and she isn't alone. She sees that these technologies could bring major advances in how schools could bring new approaches to mental health, to teaching empathy for others, and how students could learn to collaborate better with others and learn new critical thinking skills, all of which, she says, society desperately needs more of. Most importantly, she says, the teaching of science in the classroom will see huge improvements due to VR and AR.

The Grabbable Classroom

A simple rubber cube is already demonstrating that. When Franklin Lyons and Andrew Trickett started Merge Labs, which makes that cube, they had no idea that a simple cube would turn out to be their biggest seller. For a time, even after they launched the product, it looked like it might not sell well enough to support the company. They told us that Walmart had an oversupply back in 2018 and discounted the Merge Cube. They wondered if their business was done.

Luckily, school teachers swooped in and bought dozens, clearing the shelves. Since then, they have gone on to have a bona fide hit with educators and, in 2019, passed a million cubes sold.



Figure 10.2: With a Merge Cube, you can lay 3D scenes and objects on each surface and control them by turning the cube in front of a phone, camera, or AR glasses (photo credit: Robert Scoble)

The company launched in 2013 with a foam holder for mobile phones that turned them into VR headsets. Smart idea, because foam was very low cost and could take the beating that students could dish out. That led to asking the question, “What can we do with VR or AR?”.

Merge came out with a handheld controller for that early headset, but it was too expensive and too hard to make work to garner a huge audience. Plus, Oculus and HTC came out with more capable products around the same time, which took away any potential attention they could have gotten.

A better idea was about to come as they looked at the camera on the front of phones: to build a foam cube that had a unique pattern on each side of it as a sort of controllable marker system for AR. The camera sees that pattern on each face of the cube, and the phone knows which side it is looking at and how the cube is oriented in front of the phone. Then, the phone could replace the cube’s faces with all sorts of fun AR visual imagery⁵.

At first, when they released the cube at the 2017 Consumer Electronics Show, they thought it might find a market with parents who wanted new kinds of games or new kinds of stories for their kids.

Today, using the Merge Cube, you can dissect a frog or explore the universe. These new apps and new customer bases quickly went exponential. Lyons and Trickett's Merge Cube product is an early sign that huge changes are coming regarding how we educate not just our kids but also ourselves.

This new education is coming, particularly in places where classroom training just didn't fit well before, if at all. For instance, police training is best done in VR, says Steve Grubbs, CEO of VictoryXR, an education platform: "All learning is immediately in front of us wherever we are. This radically changes learning. This changes the ballgame." This new kind of training is actually a simulation of the real world that police officers move through. This makes learning far more experiential and immersive, and the learning happens much faster and sticks much deeper. Grubbs tells us there's a huge difference between being taught something, like checking a potential suspect for a weapon before moving in, and experiencing the same thing in a VR simulation.

It also is coming to homeschoolers.

Why are both cops and homeschoolers using this technology? This new kind of learning is much faster and lower cost than buying books, watching videos, or trying to video-conference into a classroom.

Homeschoolers fall into a few camps, Mark Andersen, co-founder and President of Lifeliqe, a 3D education company, tells us. One camp has one or more amazing children who need more education than regular schools, even private ones, can offer. Another camp wants strict religious schooling. Yet another camp doesn't fit into regular schooling. Maybe they don't want to comply with, say, California state laws. Maybe they had a troubled kid who was getting bullied at a regular school. Yet another camp, one that is growing, doesn't want their kids in a regular school due to the threat of gun violence. All of these camps have one thing in common, however: they are looking for better ways to educate their kids at home. In 2020, yet another camp appeared: millions of kids who all of a sudden were forced to stay home due to COVID-19. Increasingly, AR and VR will provide some major solutions.

The homeschoolers, Andersen says, are flocking to new approaches to teach their kids faster. He surmises that this might force more traditional educators to change since traditional parents will expect their kids to be taught using these techniques, too.

As homeschoolers buy more recent curriculum tools and are looking for more experiential learning to bring into their homes, they can experience the same learning other kids might be performing physically (say, frog dissections, but without the mess of the actual frog). This will provide the funding that companies like his need to turn more lessons into 3D sessions.

There's a downside, though, Grubbs reports: the customer acquisition cost of homeschoolers is higher because you have to convince them one at a time, in contrast to convincing a school administrator to buy a bunch of Merge Cubes for the entire school in one purchase.

There are lots of examples of curriculums inspiring 3D learning environments for lessons, though. The University of Alberta is teaching cell biology via VR, which presents cells that students can walk through, watch, and interact with. VR is helping children with autism learn in Bristol, UK, by dealing with their emotions and learning in a way that they take to very quickly. Mohawk College in Canada has built a \$5 million interactive digital center featuring a curriculum from EON Reality. In New Zealand, they have built an AR app to bring graphic novels to life. At a start-up, Explore Interactive, affiliated with Purdue University, students are being taught science, technology, engineering, and mathematics using a series of animated virtual cards that dance on top of tables. "Collaborative Augmented Reality is one of the most innovative things happening in education," CEO Amanda Thompson said. "It allows teachers to challenge students to work together to troubleshoot and solve problems. The potential topics are limitless."

Explore Interactive got a \$225,000 small business innovation research grant from the National Science Foundation, among other funding, and, like other companies, it is bringing students into this new way of learning. "Elementary school-age children can't work with a Bunsen burner, but with Augmented Reality, they can turn on a Bunsen burner and boil water and see how water reacts and changes," Thompson said. "Future development will allow them to zoom in to the molecular level and see the atoms heating up and moving apart from each other."

AI is playing a major role in this new learning, too. At Rensselaer Polytechnic Institute, a university based in Troy, New York, they are teaching students Chinese with a variety of AI and VR apps in a special "immersive" classroom that was partially built to study these new interactive learning environments. Inside, students talk and interact with virtual images that are actually in front of cameras that are tracking their movements and gestures as they talk with virtual beings or agents. AI analyzes whether they said the right thing at the right time in the right way. An article in MIT's Technology Review reports that the team is still in the early stages of understanding how effective it is, but in a pilot at the end of 2017, the researchers found that it qualitatively increased student engagement and enjoyment in language learning.

"It's really [moving fast] and it's hard for teachers and librarians to wrap their mind around," Lifeliqe's CEO Andersen says. But to get to the place where every student is wearing Spatial Computing devices, he says that a ton of friction has to be removed from the process. The cost is too high. Lots of wireless infrastructures that many schools don't have need to be built.

It's too hard for already overworked and underpaid teachers to figure out which apps and services could be used in their classrooms. Often, they are required to teach within the confines of tests and curriculums that aren't flexible enough to allow new approaches to be brought in. Equipment to do VR or AR, even if it's as simple as a tablet with a Merge Cube, needs to be maintained, updated, and dealt with if broken. VR headsets often can be even more confusing than traditional learning methods and paper tests since a child who is having problems can't show others what trouble they are having inside a headset. Lots of friction will make school system adoption slower than some parents might like, but we might see an effect like Apple caused with the iPhone: students might bring their own technology to school.

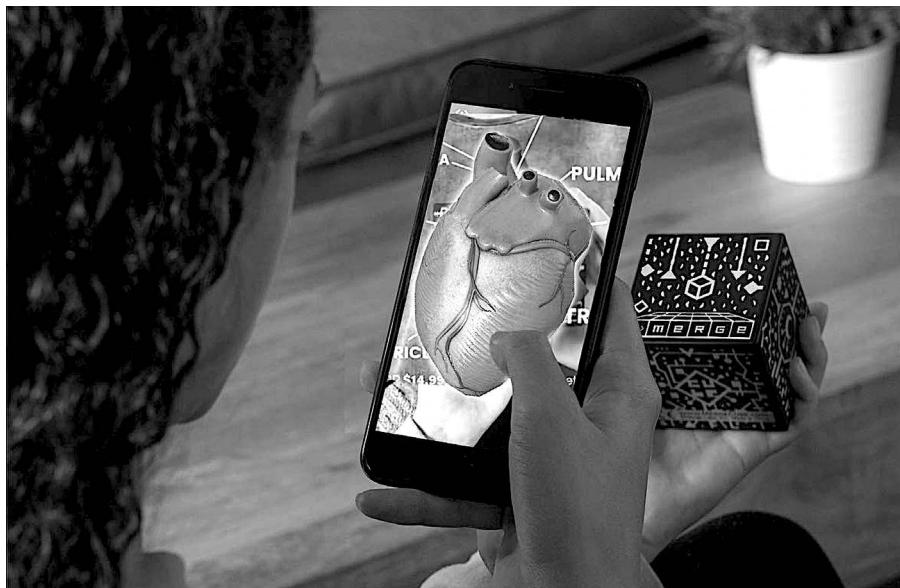


Figure 10.3: An example of what the Merge Cube can do (photo credit: Robert Scoble)

Like we saw when corporate users started bringing iPhones to work, even against the wishes of management, we predict changes will come to education whether the educational establishment wishes it or not. Parents, innovative teachers, and kids themselves will bring the best tools to school, or, if forced to, use them at home to advance their learning faster and in a more entertaining way.

Another challenge is coming for educators, though, and the needs will be bigger than they can satisfy. We all know that if trucks start driving fully autonomously, that technology will put 1.3 million drivers out of work in the US alone.

Tens of millions of people lost their jobs due to the COVID-19 pandemic and many need to be retrained for new jobs because their old jobs simply won't reappear. There aren't enough schools ready to retrain just that population. China is facing a similar problem: there simply aren't enough teachers to train the kids who are being born every year. As a result, China is building VR classrooms with AI assistants by a company called NetDragon Websoft to fill this gap in available teachers and classrooms to hold them⁶.

This may sound scary to many Western parents, but to a laid-off truck driver who needs to find a new skill to keep food on their family's table, this prospect will be a lot less intimidating. Go back to 42 Silicon Valley, the school that took a chef and helped him learn to be a programmer. Today, that school still needs an expensive building and rows of expensive computers to provide an environment for learning to take place. We predict that by the end of the 2020s, schools like that will take their curriculum, and methodology, to the virtual world, where students will meet in virtual classrooms and learn from each other—all without having to travel to a school. These new learning environments, whether in something like AltspaceVR or other virtual places that we haven't yet seen developed, will provide another nudge for the existing education system to change. Already, those changes are showing up in the workplace with new kinds of training and simulation.

Workplace Training and Simulation

When you mix computing that you, a robot, or a virtual being can move through on top of a physical machine, you get a whole new level of ability to learn a new job quickly. Whether we call this AR, Spatial Computing, Mixed Reality, or VR doesn't matter too much. It's the fact that we can get visual simulations of, say, a factory floor, and we can learn a new job as if we were on the factory floor doing the actual job or having someone standing beside us assisting us in learning that new job.

Workplace needs are going to increase, Accenture predicted in a report on Immersive Learning for the Future Workforce. It shared research from the World Economic Forum that estimated that, in 2024, nearly 42 percent of the top skills needed across all job families would change. The signals are strong that a lot of people will soon need to be retrained, and Spatial Computing offers the answer.

The same report says that learning through experience, which is what Spatial Computing brings, increases learning quality and improves retention by up to 75 percent. That number comes from studies at Stanford University and Technical University Denmark.

What else does virtual training bring? Well, much better analytics so that you can prove your workers were trained, for one. Our Spatial Computing devices already capture a ton of data from eye, motion, behavioral, and gesture sensors, and that will all be used to ensure that training is being understood and can even make sure that the job is performed properly on factory lines, in mines, and in many other jobs. With VR and AR, there is the additional benefit of being able to be educated about dangerous tasks in a safe indoor environment.

It also lets you learn without fear of the consequences of mistakes. At Boston Scientific, they showed how doctors can practice heart surgery over and over again without fear of messing up and costing a patient their lives. VR and AR provide real-time 3D views of human anatomy that can be significantly more effective than how we learn from diagrams or 2D training videos. Immersive learning experiences like this are also what today's tech-savvy generation of students expect from healthcare education. Other companies that are using Spatial Computing for healthcare, surgical learning, and simulation include Dr. Shafi Ahmed's company, Medical Realities, Medivis, and Osso VR.

What is wonderful is that the benefits of using Spatial Computing for learning are applicable to workers doing far more mundane jobs like putting an air conditioning unit into a car or a toilet in a bathroom. Why do we bring that up? Because Accenture reports that workers who had to learn to do that task completed it 12 percent more accurately and 17 percent faster after doing training in VR rather than watching similar training on a 2D video.

We have dozens of such stories of companies using VR to enhance training for these mundane tasks, but something deeper is underway. As we talked to people who are doing VR training, we kept hearing about training that was hard to do before VR came along.

If you visit Ford's plant in Detroit where they build F-150s, as we did, you will see a new one is built every minute or two—24 hours a day. They rarely stop the line. After all, if you owned a system that spits out \$40,000 every minute, would you ever stop it? This shows why so many factories are already using VR to train their workers somewhere away from the factory to do difficult and often dangerous jobs. In VR, you can slow down the line virtually. If you make a mistake, there aren't consequences like there are on the real line.

You hear factory workers learning new terminology because of this shift. "Digital twin" is one you frequently hear. This is a virtual version of the entire factory floor, one you can experience in either a VR or AR headset.

Training isn't the only thing that's changing on manufacturing lines, and we cover more of them in *Chapter 6, Vision Three – Augmented Manufacturing*, about how manufacturing and warehouse work will radically change over the next decade, but it's worth digging into workplace training here. It's not the only change happening to how workers are learning new jobs either.

Another factor at play here is that older workers are retiring at a fast rate from many jobs, often leaving factories with decades of knowledge. Using Spatial Computing, they can leave lots of that learning on the factory floor by recording videos and leaving those videos on top of the digital twins that are now being built. That's helpful because, mixed with other types of training, you can find that wisdom more easily than using a search engine and it's on top of the machine you are being trained to run.

We see many workers wearing devices like RealWear's devices, which don't really present much of an augmented view of the world; rather, these devices mostly augment work by presenting information on small displays, even video, and open up new kinds of remote training capabilities. These devices augment jobs in real time, helping workers do things around factories, oil refineries, and other enterprises, and they also let you learn remotely from experts elsewhere in the enterprise. Imagine you are installing air conditioning units on, say, that F-150 truck, and you hit an issue you never have seen before. Quick! You can't stop the line for long since every minute of stopped time costs your employer at least \$40,000. Do you reach for a paper manual? Or do you call your boss or another expert to ask for help? That's exactly what's happening now. Microphones on the RealWear device let you talk and a camera on the device lets the remote expert get a look at the problem via video conference. There's a screen on the device that lets you see any visual solution to the new problem, too. This is all in a head-worn device that lets you continue to work with both of your hands.

Workplaces are changing already because of Spatial Computing technologies. Today, we are designing things in 3D, like new kinds of architecture and new kinds of products, yet many of us are seeing those things on 2D screens, which really don't map to how our brains work very well and it's hard to visualize how those products will be used in real life. Architects took to VR very quickly. We were shown one home completely in VR and we could walk through it and even visit the pool outside and see how it would look in the real physical space where it is being built. (The home is being built for Lynda Weinman and her husband. She started Lynda.com, which was one of a group of companies that radically changed learning thanks to the internet.)

Others have to learn new kinds of jobs very quickly or deal with engineering changes to existing products or lines. Other kinds of jobs, like firefighting, criminal justice, or police work, are changing every day because of new science or new technology, and here, too, VR and AR are already seeing radical new uses. The New York City Police Department is using VR to train their police on how to act during high-stress events, such as terrorist, hostage, and mental health emergency situations.

That brings us to the role of simulation in workplace training. Some things happen so infrequently, are so dangerous to train for, or are so costly to recreate in real life, that real-life training is, for the most part, intractable. Learning to defuse a bomb, for instance, is a bit hard to do on a real bomb, yet we met people who had to learn how to do that when we visited Nellis Airforce Base in Nevada, where we were taught how to defuse such a bomb. We imagine that learning that skill in VR is a bit more comfortable than having to learn on a real device. Even if you say learning on a real device is better, it's hard to teach a large number of people quickly about a new kind of device. Virtual training can be passed to your entire workforce immediately, which can save lives and money and ensure that everyone is on the same page. Training like this can also be transmitted to soldiers in the field via Spatial Computing systems, and when one base discovers a new kind of bomb, that training can be sent to teams all over the globe within minutes (see *Figure 10.5*).

These are similar changes to those that will roil the previously stable waters of classrooms over the next decade. These changes are already starting to happen but will greatly increase as Spatial Computing devices that consumers and workers can afford and that are light enough to wear for hours every day get released by companies like Apple, Facebook, Microsoft, and Sony. Already, many are spending hours every day inside Meta Quest 3 headsets, which cost \$500, something that many workers can afford.



Figure 10.4: A surgeon practices, in VR, a new surgical technique using Osso VR's surgical preparation system (photo credit: Osso VR)

Why will this happen? As we already explained, humans tend to remember things better when they are experienced more directly, as they do when experiencing VR or AR. Because of this, less time is spent on learning and more can be spent on the job. Strivr says that VR training improves employee confidence and has boosted test scores by 10 to 15 percent, even for trainees who simply watched others experience the training.

Additionally, we are seeing a move to more jobs being done remotely, even those who coordinate with people on factory floors. With the use of VR and AR, this will only increase. Younger workers are hoping to work this way, and if factories can keep some people working at other locations or working using VR and AR headsets at home, traffic and carbon footprint problems could be alleviated.

Another reason for the increasing adoption of VR and AR for workplace uses is that you can experience things from various points of view. For instance, you can become a woman experiencing sexual harassment, or the boss that is dishing it out. This increases empathy, which could be an HR training objective, and the same hardware could be used for factory training so you can walk around a machine and see how you would work on it from any angle.



Figure 10.5: A defused bomb for training purposes at Nellis Air Force Base with an X-ray of the scan. (photo credit: Robert Scoble)

At BMW, workers use VR to practice their jobs before they get onto the factory floor. They can grab virtual parts, see how to install them, and practice doing that with virtual tools, and after that is all done, the system gives them a score on how well they did each task. Armin Hoffmann, project manager at BMW's group plant in Landeshut, says that this lets his teams learn how to build new cars or get updated on new part designs faster, which saves BMW money and reduces error rates on its factory floors.

Verizon uses VR to train employees on what to do during a robbery in its stores. The New York Mets use VR to practice hitting against tomorrow's pitchers in hotel rooms the night before games. Norway's World Cup soccer team used it to prepare against its opponents. In Strivr, you can practice firing an employee. In Ireland, the military practices sweeping for mines in VR. In some medical centers, you can learn how to deliver a baby. We have hundreds of examples of training and simulations already being used on many other jobs.

VR training isn't just replacing previous methods of training, although, in some cases, it is, like with sexual harassment training, where VR can bring new approaches to getting into someone else's shoes so you can feel how your words and actions might hurt. It is also excelling in new, more dangerous areas, where previously, the best thing you could do was read a book about it.

Walmart, for instance, takes new hires onto the retail floor during Black Friday, which is its busiest day of the year. You can watch as many videos as you want or read entire books about what it's like to deal with customers who are fighting to get a deal on a TV, but VR can actually take you there and give you a much more real experience.

Shelley Peterson, who was Lockheed Martin's emerging technologies lead and principal investigator for AR, reported that their efforts to get training out of the "rectangles" that represent paper or screens that most of us use are already paying off big time: "It has a significant impact. You can capture a concept much more thoroughly when it's not in that rectangle of data. It changes when you are picking up a new concept."

She explained that often, when building spacecraft, workers have to crawl through tight, confined spaces, sometimes with potential hazards, and other workers have to work in clean rooms that take minutes to undergo sterilization procedures and suit up so as to meet requirements. Even then, workers going into such areas can introduce defects to the manufacturing process, and other workers have to crawl up and down lots of ladders. If a class is following a worker to help them learn a new task, that alone can take minutes away from learning.

AR makes learning much better because it shows the students all sorts of things that they need to be aware of, and virtual interfaces make many physical tasks easier. Then, AR headsets (in Lockheed's case, Microsoft HoloLens 2) show whether or not you did the job properly.

"Where it normally would take 8 hours (to comprehend), now it takes about 45 minutes," said Shelley. This saving of time is there even with seasoned technicians, she reports, while pointing out that AR even saves time when building the instructions, and that they saw gains in cost, quality, and time to get the job done.

Lockheed isn't alone. Research by Capterra in 2021 found that almost 40 percent of small and midsize US businesses were expecting to use Spatial Computing in employee training by 2024 and projected that it will get employees up to full productivity 50 percent faster than traditional instruction. By comparison, in early 2018, 32 percent of retail employees said they didn't receive any formal instruction, which shows there's a big potential market for Spatial Computing technology to be used here.

The PR consequences for not doing the training are extreme now. Sephora closed all its doors for an hour to perform "inclusivity workshops" following accusations of racial profiling. Starbucks similarly captured headlines when it closed 8,000 stores for an afternoon of racial-sensitivity training after two African-American men were arrested at one of its stores while waiting for a friend, *Forbes* reported.

The list goes on and on. Lowe's is using it to help employees learn about how to better perform do-it-yourself projects. Chipotle is using VR to help with food safety. In VR, you can learn how to drive a bus, be a railway engineer, or get chemical industry training. At Airbus, they are currently using HoloLens headsets to virtually test their designs to see if they are ready for manufacture. It reports that using Spatial Computing speeds up the process substantially, decreasing the time spent in training by 80 percent. Another VR use case is for lifeguard training for the American Red Cross, which was rolled out to over 300,000 lifeguards across 12,000 training centers in the United States⁷.

We've seen this ourselves. In Bailenson's lab at Stanford University, his students wanted to teach people about climate change and have the facts "stick." We wore a VR headset and were taken underwater where we experienced firsthand the damage being done to the oceans. It showed us how carbon dioxide molecules, when absorbed into the sea, created acidic conditions and were killing the reefs. That memory sticks clearly in our minds even two years later.

Speaking of climate change, so much corporate training requires students to fly long distances to training rooms. Just flying a short distance from San Francisco to Los Angeles for a one-day training costs \$300 for the flight, \$250 for the hotel, and \$75 for a meal. That's already more than a single Meta Quest 3, with which you can experience the same training in VR, and that's not counting the value of the worker's time spent traveling. If you have to fly your workers internationally, the cost savings become even more extreme. With VR, you can teach everyone remotely at the same time.

We saw a class for real estate agents that was done completely in VR. No travel required. More pressure will be put on companies to do this kind of training due to the carbon costs of flying workers around the world as climate change becomes a much more serious issue for both customers and investors. As a benefit, your workers might get together in Rec Room and play some games, too, which will increase the bonding your team has without the problems that come with tons of alcohol consumption in the hotel bar.

Other benefits? The “teacher” never has a bad day or forgets to communicate something important. It’s there every time for every student and, due to the sensors that watch where a student is looking and virtual tests that can be done, you’ll have proof that the job was understood and remembered at a far higher rate than if you don’t have VR and/or AR training.

With trades such as healthcare, military, and police that could inherently provide safety risks, VR provides a safe training alternative. For example, the New York City Police Department is using such simulations to train for new kinds of terror attacks. This technology is rapidly developing, thanks to not just the better visualizations that are possible in Spatial Computing headsets but also the AI that can run all sorts of things. We learned about this by studying autonomous cars. Car companies like GM and VW are running billions of miles where a virtual car drives around virtual streets that are so realistic now that they even have virtual puddles due to recent rain. That same technology is currently being used in workplaces.

Commenting on VR workplace training that uses simulation, Andy Trainor, former senior director of Walmart Academies, said the following: “We can create scenarios that may never happen (like severe weather training) or situations that happen all the time, but we don’t want to disrupt our stores to give associates firsthand experience. VR lets them explore these sorts of scenarios in a firsthand way.”

The workplace location where Spatial Computing could be used for simulation could be anywhere, including a factory, logistics facility, retail storefront, or even US military and Air Force training facilities and battlefields. HoloLens helps Airbus aircraft designers virtually test their designs to see if they are ready for manufacture.

The use of the HoloLens sped up the designing process prior to manufacture substantially, decreasing the time spent by 80 percent. In addition to this, Airbus has identified more than 300 use cases for AR and they will be continuing to use the HoloLens 2 headset.

Given that there are close to 35,000 serious injuries involving forklift accidents each year in the United States alone, the Raymond Corporation, which is a leading manufacturer of forklift trucks and pallet jacks, created a VR simulation for forklift operator trainees. 3M, Caterpillar, and UPS are other companies using VR simulations to enable workers to know how to operate in situations that are potentially not safe.

This brings us to a new kind of training that was hard to do before: teaching humans to manage and behave better, both while working with other humans and while working with other machines. It turns out that early science shows that if you have an emotional connection to the robot working alongside you, you will perform better yourself.



Figure 10.6: A worker uses VR to both build and train on a new manufacturing line (photo credit: BMW)

A Virtual Arm for Human Resources

As anyone who has worked in a corporate organization can tell you, one of the first stops when you get a new job is the company's HR department. There is a terrific amount of work that an HR department does for a company, but it usually flies under the radar of most workers since its touchpoints are seemingly on the periphery of what is considered to be the main company business.

However, if a company's HR department is lax and careless, workers will inevitably suffer—from minor issues such as losing an accounting of a vacation day or two to losing large sums of money due to sloppy equity options paperwork.

A company's culture is usually governed by its HR department. Empathy, diversity, and sexual harassment training come under its purview. Official workplace safety is another area. If there are problems in these areas, upper management and HR are the first entities to be blamed.

Additionally, a company's HR department has a large part to play in developing and enforcing the rules and standards for pay grades, reviews, and promotions.

In all of these areas, VR has already been implemented, sometimes as a complete replacement, and other times as an adjunct, for a company's HR tasks. Examples of tasks covered by VR include recruiting and onboarding, benefits training, empathy and diversity training, sexual harassment training, workplace safety, and guidance on reviews and recommendations for promotions.

Regarding workplace safety through the HR department, over 11,000 Intel employees and contractors are required to take an electrical recertification course every year—this course is now given in VR.

Talespin by Cornerstone offers VR workplace simulations so that employees at all levels of a company hierarchy can practice and learn how to deal with others in workplace situations.

A company that has taken the VR role in HR further than most is Walmart, with its training vendor Strivr. Walmart is using VR to determine their workers' skills and overall potential on the job using VR simulations. The data gathered during the simulations is then used as part of the decision-making process to determine who will get a promotion or a pay cut. Currently, data gathered is based on how workers answer questions put to them while in the VR simulation. However, plans have been made to gather workers' body movement and attention data while in VR, which early research has shown provides a relatively more accurate picture of future on-the-job performance.

Our Learning Futures

Thanks to Spatial Computing technologies, we'll see a deep change in how we educate both our kids and ourselves. Soon, we'll be learning on top of tasks or via simulations. Our learning will be far more immersive and interactive than anything we've experienced before, and we predict, eventually, far faster and at a lower cost, too.

As we learn, we can teach, too. The AI systems that will study our every move can learn from us and even start predicting what we'll do next. That will further take our learning to a much more efficient place because soon we'll be learning while doing. Imagine a pair of glasses that notice you are about to touch a hot stove. "Hey, that's going to burn," our systems might say just before we put our hands on that hot surface. Think that's impossible? Not for computers that are starting to do a variety of things from driving cars to helping us shop with these predictive systems. Let's go into that next.

References

1. MoneySense (mobile-only): https://ulsterbankni.mymoneysense.com/media/8590/5_8_ar_games_parents.pdf
2. Stem Arts Lab: <https://www.stemarts.com/>
3. Strivr Uses Virtual Reality To Train #NFL Quarterbacks: <https://www.youtube.com/watch?v=XwnjjUDOets>
4. Berger L.M., Wood G., Kober S.E. Effects of Virtual Reality-based feedback on neurofeedback training performance -A sham-controlled study. *Front Hum Neurosci.* 2022 Aug 12;16:952261. doi: 10.3389/fnhum.2022.952261. PMID: 36034118; PMCID: PMC9411512. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9411512/>
5. How does the Merge Cube work? <https://www.youtube.com/watch?v=aQ-yb4chHWk>
6. NetDragon Introduction Video: <https://www.youtube.com/watch?v=bvk6EhYoABc>
7. Lifeguard VR: <https://presentation.lucidrealitylabs.com/VRLifeguardsTrainingInnovation/>

Part III

The Spatial Business

Productivity, Efficiency, and Making Profits Are About to Undergo Radical Changes

This section comprises the following chapters:

- *Chapter 11, The Always Predicted World*
- *Chapter 12, The Business of Generative AI*
- *Chapter 13, The Impact of the Apple Vision Pro*
- *Chapter 14, How Human?*

11

The Always Predicted World

Humans are predictable. If we go to church, we are there at the same time every weekend. We head to work at the same time. Kids have to be dropped off at the same time, to the same place. Even things that seem somewhat random, like going to grocery stores or for dinner at restaurants, have a pattern to them.

What is your next best action? Humans are pretty good at figuring that out, right? But computers can assist you even more by warning you of, say, an accident ahead on your route, or some other pattern you can't see ahead of time: "Ahead, there's a new pothole in the left lane."

Here, we'll explore the predictive capabilities that will come with Spatial Computing and how these technologies will change almost every part of our world, from transportation to education.

The Predictive Breakthrough

Predictability goes beyond knowing that if it's 8:45 a.m. on a Monday, you probably are headed to work. Vic Gundotra, who used to be an executive at Google, told us that Google realized that if you walked into a store, that act predicted you would end up buying something in that store. That might seem obvious, but it was a breakthrough in thinking through how to advertise to users. Kelley Blue Book, among others, used that learning to put car advertising in front of you once you walked onto a car lot, via your mobile phone's notifications or by sending you an email based on your previous login to Kelley Blue Book, sometimes even convincing you to leave that lot to get a better deal elsewhere.

If you are visiting a Volvo lot after doing a bunch of searches on Kelley Blue Book, another car company might even offer you \$50 to come to their dealership and take a test drive in an attempt to get you to leave the Volvo lot without making a purchase.

As Spatial Computing evolves, these predictive capabilities will become even more advanced. The ability to anticipate user behavior will impact virtually every aspect of life—education will be tailored to students' learning patterns, urban planning will respond dynamically to traffic flows, and personal healthcare will become increasingly proactive based on daily habits and environmental data. The potential to transform everything from retail to public infrastructure is immense, and Spatial Computing is the key to unlocking a smarter, more predictive world.

Glasses That Understand You

Those early uses of contextual data and predictive systems will soon seem pretty simple when compared to what's coming next. That's a world where many users will wear Spatial Computing glasses, with tons of cameras and sensors to watch what the wearer is looking at, talking about, and interacting with. This world will arrive at about the same time as autonomous vehicles and delivery robots start showing up on our city streets. Each will have a ton of sensors as well, gathering data about the street and sharing it with a new type of cloud computing that can learn exponentially and deliver radically new services and even virtual assistants, all due to massive amounts of data being collected about the real world and everything and everybody interacting in it.

The software of Spatial Computing glasses will receive a very good understanding of the context that a user is in. We call the AI engines that will watch all of this "exponential learning systems." The AI will continuously learn about your habits and preferences, constantly improving upon its former knowledge. After all, you will want to do different things on computers if you are shopping from if you are trying to work in an office, and those two things are quite different from if you are in a restaurant on a date night.

These glasses we soon will wear will have three new systems: data collection, contextual deciders, and exponential learning systems, which will constantly use data and user context, along with other predictors such as who you are with, with the aim of making your life better. These systems have been under design for at least 20 years but are now getting much more complex due to new technology, particularly on autonomous vehicles, which have AI that does all sorts of things, including predicting whether a pedestrian at the intersection ahead will cross the road or not.

At Microsoft Research, back in 2005, researchers showed us how they built a system that could predict traffic patterns in Seattle 24 hours a day. It knew, for instance, that it was pretty likely that you would get stuck in traffic coming across the I-90 bridge during rush hour. This might seem pretty obvious to most humans, but it was an important breakthrough for computers to be able to predict traffic conditions before they happened. With Spatial Computing glasses, similar AI-based systems might make thousands of similar predictions for their wearer.

These predictions, and the assistance they will give, will change human life pretty significantly, and provide the ability to build new assistants that go even further than the human assistants that executives use today.

For instance, imagine walking into a crowded city intersection. Your glasses might instantly analyze the scene, predicting whether a pedestrian ahead will step into the crosswalk or notifying you of a potential obstacle. In a retail context, the glasses could recommend special offers as you browse a store or even suggest alternative products based on your preferences. At work, they could organize your digital workspace based on tasks you typically perform, making the experience of managing information faster and more intuitive.

The impact on everyday life will be profound. The predictive power of these glasses will introduce virtual assistants that provide personalized support far beyond what is currently offered. Unlike today's executive assistants, these systems will handle countless decisions and tasks—anticipating needs, simplifying workflows, and even suggesting optimal solutions to complex problems. The assistance provided will change the way we interact with the world, making daily life more streamlined and efficient, and opening up entirely new ways to approach both work and personal life.

As these glasses evolve, their applications will go well beyond novelty or convenience. They will empower us to navigate a rapidly changing world, helping us understand complex environments and the probability of better decision-making, all through the seamless blending of the digital and physical realms.

(Human) Assistance Not Required

To understand how these changes will occur, we talked with someone who is an experienced executive assistant. She doesn't want to be quoted or have her name used, but she was an executive assistant to Steve Jobs and still is an assistant to an executive at Meta. Executive assistants do everything from getting other executives on the phone to doing complicated travel planning to maybe something simple like ordering lunch for the team. She told us stories about getting the CEO of Sony on the phone and helping to fix problems for Steve Jobs. In other words, she cleaned up messes, since Steve was famous for getting angry when something didn't go right. She has a deep Rolodex of other executives' cell phones and their personal assistants' details that she could use to make such meetings and other things happen.

The thing is, her kind of job is quickly going away. The Wall Street Journal reported that more than 40 percent of executive assistant positions have disappeared since the year 2000. Why? Well, think about your own life. In the 1970s, travel, for instance, was so complex that you probably would have used a travel agent to do it.

Today, mobile phones and services like Google Flights have made that job much simpler, so you can do travel planning yourself. Heck, you can even get a ride to the airport by simply using a mobile app and you know if that ride is on its way, which wasn't possible just a decade ago. The same is happening to all parts of the executive assistant's job; after all, now that you can say "Hey Siri, call Jane Smith," you probably don't need an assistant like our source to get someone on the phone.

Soon, virtual assistants in our glasses will have a lot more data to use to help us than just looking at the number of steps we are taking or our heartbeat. We've seen AI-based systems that know how many calories you are consuming, just by taking a few images of the food you are looking at, and Apple has patents for its future Spatial Computing glasses that watch a number of different health factors, including listening to your voice, heart rate, and other biometric data to predict how your workout is going. Apple has a set of patents for a bunch of sensing capabilities for looking into the human eye for a variety of ailments, too.

These aren't alone—Robert Adams, the founder of Global e-dentity™, showed us how his patent that looks at your vascular and bone structure could be used for a number of things, from verifying your identity to assessing your health. Future glasses will have cameras aimed at different parts of your face for just these kinds of things. Our virtual assistants will have a number of helpful features made possible through Spatial Computing, which could allow them to see to most, if not all, of our assistance needs.

These AI-powered systems won't stop at health monitoring. The data collected by Spatial Computing glasses and integrated AI will revolutionize how we receive assistance in all areas of life. Virtual assistants will have a profound understanding of our routines, preferences, and needs. Imagine glasses that predict when you're about to run late for an important meeting, or that automatically adjust your schedule based on traffic conditions—all without you lifting a finger.

These assistants will continuously learn and improve, becoming more adept at handling complex tasks that were previously the domain of human assistants. Whether it's managing personal schedules, recommending the best route based on real-time data, or even interacting with people on your behalf, AI-powered virtual assistants will soon handle it all.

In short, the future of assistance is about more than convenience—it's about predictive intelligence. Virtual assistants, powered by Spatial Computing, will transform how we interact with the world and with the technology around us.

Pervasive

Now, imagine the other data that such a system would have if you are almost permanently wearing a camera on your face—developers could build systems to play games with you about your eating habits that would say to you, “Hey, don’t eat that donut, but if you do, we will remove 400 points from your health score and that will require you to eat 10 salads and go on 2 extra walks this week.”

Already, we are seeing many computer vision apps that are starting to identify every object you aim the camera at. They are using the same kind of computer vision that a self-driving car uses to know it is seeing a stoplight, stop sign, or another object on the road (while writing this chapter, Robert’s Tesla started recognizing garbage cans, of all things).

If you use Amazon’s app on a mobile phone, there’s a camera feature that does something similar. Aim the camera at a Starbucks cup and it will show you things you can buy from Starbucks. Aim it at someone’s black-collared shirt and it will show you other collared shirts. In the Spatial Computing world, you will see these capabilities really come to bear in new kinds of shopping experiences that can take a single 3D object, like a scan of a coffee table, and show you lots of other things that will fit with that table.

What if, though, it knew you needed a shirt before you even started looking for one? How would that work? Well, if you put a new job interview on your calendar, a predictive system might surmise that you will need a new suit to look spiffy for that interview. Or, if you learn you are pregnant, such a system might know that you soon will need a new car, new insurance, and a new home. In fact, such predictive systems are already running all over the place and are the reasons why after a big life event, we get a ton of junk mail advertising just these kinds of things. The New York Times reported that a system at Target knew a woman was pregnant before she told her family because she started changing the products she was buying and a computer system recognized the change in her patterns.

We can see these systems at work in tons of other places, too, including inside email systems, like Gmail, that start answering our emails with us and giving us predetermined answers that often are pretty darn accurate.

The world we care about, though, isn’t the one of email; it’s the new computing world that’s 3D and on our faces. Here, predictive systems will be watching our eyes, our hands, our movements, and everything we touch and consume. Now, don’t get all *Black Mirror* and dystopian on us.

Truth is, the early devices won't do that much due to the small batteries that will be in them and the sizeable privacy concerns, but over the next decade, these devices will do more and more, and, when joined with devices in our homes, like Amazon Echo or Google Home, will start doing a sizeable number of new things, including changing the way we go to the store, using what we call "automatic shopping." We'll have versions of that store virtually in our living rooms.

But let's take this even further. Predictive systems won't just stop at health warnings or shopping suggestions. Imagine that, while you're walking around your home, your Spatial Computing device anticipates your needs and proactively offers assistance. Did you leave the lights on in another room? The glasses will remind you. Has it been a while since you had water? The system can recommend a hydration break based on biometric data, such as skin temperature or heart rate. This level of assistance, driven by predictive intelligence, will become the norm, anticipating your needs before you consciously recognize them.

In education, predictive systems could track your study habits and optimize learning plans for you. If you're losing focus or struggling to understand a certain topic, the glasses could adjust the lesson format or provide real-time feedback to keep you engaged.

In entertainment, you might be watching a movie at home and suddenly decide you want to recreate the feel of a cinema. With smart home integration, the predictive system could dim the lights, adjust the sound, and even recommend snacks—anticipating your preferences based on prior movie-watching experiences.

In travel, the predictive system could assist with packing or route planning. Based on your calendar or travel history, the system could remind you to pack items you tend to forget, like a toothbrush or charger. And, as you approach the airport, it could alert you to nearby traffic or potential delays, recalibrating your route and ensuring you arrive on time.

When it comes to automatic shopping, predictive systems can take convenience to the next level. Imagine a future where you never run out of basic household items. Your Spatial Computing device will monitor your consumption patterns—how quickly you go through groceries, toiletries, or even your favorite snacks. When you're running low, the system could automatically reorder these products, ensuring they arrive just as you need them, without you even having to think about it. This could extend to bigger purchases, like furniture or clothing, with the system providing timely suggestions based on wear and tear or style trends.

You might be getting ready for a special event, like a wedding or a holiday, and the predictive system could anticipate your needs weeks in advance. It could offer recommendations for attire, gifts, or even travel arrangements, all curated to your preferences. It could suggest booking a hair appointment or ensuring your wardrobe is up to date, sparing you the last-minute rush.

As Spatial Computing becomes more integrated with devices like smart mirrors, wearables, and even vehicles, the predictive systems will work across platforms, providing a unified and seamless experience. You could check your health stats while brushing your teeth in the morning, get shopping suggestions while cooking breakfast, and receive fitness recommendations while driving to work, all through interconnected devices.

The world of predictive systems will expand far beyond simple notifications or suggestions. These systems will anticipate the subtleties of our daily lives and adapt to our evolving needs. Privacy concerns will need to be addressed as this technology becomes more pervasive, but the potential for improving daily efficiency, personalization, and convenience is vast.

We are entering an era where Spatial Computing will not just augment our environment but also fundamentally shape how we live, work, and interact with the world around us.

Proactive

Going back to the human assistant, if you really had an assistant who knew you well, especially if you had one who lived with you, that assistant would know if you were running out of milk or cereal. Well, already, making a shopping list is easier than it used to be. In our homes, we just say “Alexa, add Cheerios to my shopping list” and it does. But within a few years, the glasses you are wearing will just make note that you have had your seventh bowl of Cheerios this month and that there are only two bowls left in the box in your kitchen, so it will just add another box to your shopping list. The AI driving your glasses will also be able to predict what you might want to order and ask if you want those items. Yeah, you can take it off if it gets too aggressive, but we bet that most of the time, it’ll be pretty accurate. Added onto new delivery robots, you’ll see almost a total change in how we stock our kitchens.

When we were writing the first edition of this book, our Tesla started doing something similarly automatic: “automatic navigation.” Instead of having to put an address into the car’s navigation system for, say, heading to a doctor’s appointment, the car looks at your Google Calendar and then automatically navigates to your next appointment when you simply get into the car. Hey, it’s pretty accurate. After all, you did add that appointment to your calendar. Where it wasn’t accurate or didn’t have enough data to work, it trained us to make sure our calendars were accurately kept up to date with the street addresses of appointments.

This two-way feedback system proves very effective, and its accuracy increases our trust in, and love of, our car. Other owners have reported the same. One thought it was the most amazing new feature because it saved her from trying to find addresses while driving or touching the screen while her nails were drying on the way to work.

That's one level of prediction—a system that just does pretty simple stuff—but what will really change radically over the next decade is when you hook up the exponential learning systems of AI and computer vision, along with millions of cameras and sensors moving around the real world.

In the near future, these proactive systems won't stop at merely helping with shopping lists or calendar reminders. As AI continues to advance, the systems in your Spatial Computing glasses will take this proactivity to entirely new levels. Imagine your glasses identifying subtle changes in your health, like your posture, breathing patterns, or even your tone of voice, and offering personalized recommendations to improve your well-being.

For example, if the system notices that you've been sitting at your desk for too long, it might suggest taking a short walk to stretch your legs and boost your energy. Or, if it detects elevated stress in your voice during a meeting, it could recommend a quick mindfulness exercise afterward to help you reset. These systems will know when to step in and provide assistance based on your context, without you needing to prompt them.

In work environments, these predictive systems will become crucial productivity tools. Your glasses could automatically organize your workspace based on what you're currently working on, pulling up the documents or software you need without you even asking. During meetings, the system might proactively take notes, capturing key points and summarizing them for you by the end of the discussion, or even remind you of follow-up tasks you committed to during the conversation.

In education, predictive systems will revolutionize learning. If you're studying a challenging topic, the system could detect when you're getting stuck and provide additional resources or suggest alternative learning strategies based on how you've successfully grasped information in the past.

Imagine waking up and, before you even leave your bed, your Spatial Computing glasses have already begun adjusting your environment to suit your daily needs. Based on your sleep quality, the weather outside, and your calendar, the system could recommend an outfit and suggest breakfast options to match your energy levels. If you have a workout planned later in the day, it could remind you to hydrate and keep track of your activity levels throughout the day.

Beyond the simple tasks of adding items to your shopping list, these proactive systems will integrate seamlessly into your daily life, offering personalized assistance at every step. As you run low on pantry staples, the system could automatically reorder items or, based on your dietary habits, it could suggest healthier alternatives to your typical groceries. This won't just apply to basic goods either—whether it's reminding you to renew your car insurance, suggesting a new route for your morning jog, or adjusting the thermostat to your preferred temperature when you arrive home, these systems will be finely tuned to your personal preferences and routines.

Proactive systems will extend beyond individuals and into smart homes and cities. Your smart home will adjust its environment based on predictive data gathered throughout the day. If you tend to lower the blinds every afternoon to avoid glare, the system could learn this behavior and automate it, adjusting the lighting and ambiance based on your schedule. If you're on your way home from work, your smart home could preemptively set your preferred temperature and lighting levels, even selecting a playlist for your evening relaxation, all without needing a single voice command.

In smart cities, these predictive systems will guide public transportation, traffic management, and even waste disposal based on real-time data. Sensors across the city will work in harmony with personal devices, optimizing everything from pedestrian crossings to energy usage, ensuring that the urban environment adapts to the flow of daily life.

Proactive systems will revolutionize healthcare as well. If your Spatial Computing glasses notice abnormal biometrics—like a change in your heart rate or skin temperature—they could suggest scheduling a check-up or provide on-the-spot advice. These systems could even assist healthcare professionals by collecting and analyzing data continuously, allowing more accurate diagnoses and personalized treatment plans. If a patient regularly experiences fluctuations in their glucose levels, for example, the system could notify them before their levels become problematic, offering suggestions to stabilize their health, such as consuming a snack or taking a short walk.

In emergency situations, these proactive systems could be life-saving. Imagine your Spatial Computing glasses detecting an irregular heartbeat and automatically alerting emergency services, providing them with your exact location and medical history. By predicting potential health crises before they happen, these systems will empower both individuals and healthcare providers to prevent rather than react to problems.

The power of proactive systems lies in their ability to anticipate your needs and enhance your day-to-day experiences, all while learning and improving over time. From simple tasks like reminding you to pick up groceries to complex actions like predicting health concerns or optimizing entire city infrastructures, AI and Spatial Computing will work together to make your life smoother, safer, and more personalized.

As these systems evolve and learn from your patterns, they will become an invisible yet integral part of your daily routine—offering guidance, improving efficiency, and, ultimately, enhancing the quality of life for individuals and communities alike. We're on the brink of a world where everything around us, from our homes to our cities, will work intelligently and proactively to meet our needs, often before we even know we have them.

Connected

What do we mean by “exponential learning system”? Well, one night, when walking home from band practice, a fellow student back in 1982 was hit in front of the school and died. The next year, a stoplight was erected in front of Prospect High School in Saratoga and it is still there 35 years later. That system was a linear learning system. It improved that one intersection. If you lived in, say, France, though, it didn’t improve your life at all.

Now, compare that to how an original June Oven got its information and made appropriate changes. That June Oven had a camera inside and an NVIDIA card that let it do machine learning. Why would you want an oven like that? Well, if you had put a piece of fish into the June Oven, it said “salmon” and correctly cooked it without you needing to look up the recipe. What if, though, you put something into the oven it couldn’t recognize, such as Persian kabobs? Well, then you could set it manually, which taught the oven about something new. A photo of that kabob, along with the temperature and time settings, were sent up to June’s servers, where programmers there could then add that to everyone’s oven, which made everyone’s oven exponentially better.

It’s the same for Tesla cars. When we slam on our brakes in our Tesla, it uploads the video from that part of our drive, along with all sensor readings. Programmers there can then work that into new machine learning training, improving the system for, say, understanding what a pedestrian looks like, and improving the system for the entire fleet on the next update. These new products, which have neural networks inside, learn exponentially and get dramatically better after you purchase them, and do it fleetwide. So, unlike the stoplight example, which only improved one intersection, the next time something goes wrong, the programmers will fix the program for the entire fleet. It gets better exponentially and, as more users get these products, the time to new features goes down exponentially as well. Within a decade, these exponential learning systems will dramatically change almost every part of our lives, particularly inside Spatial Computing glasses, which have an exponential learning system built in.

Back to Spatial Computing glasses, then. When you first get them, they might not recognize much in the world, but if millions of people have them and teach them one new thing a week, that means your glasses will recognize millions of things within a few weeks, exponentially learning and getting better with a rate that will speed up as more people get the glasses.

What this means is that we will soon have virtual assistants and coaches who get smarter and smarter every week. Our virtual running coach might start out only understanding simple runs, but within a few weeks, might be taking us on back-country trails that are poorly marked.

Our virtual nutrition coaches might start out only understanding major brands but soon will recognize unusual candies from far-off places. Our shopping services might start out not knowing our favorite brands, but within a few weeks, will learn exactly what brands and colors make us happiest.

Our Spatial Computing glasses will bring this kind of new utility, along with some that might even save your life and take the fall detection feature of the Apple Watch family way further. Let's dig into just how deeply human life will change because of this data collection, contextual decision-making, and processing, and the exponential learning systems that Spatial Computing will bring.

Beyond personal coaching and daily convenience, exponential learning systems will evolve into more complex and specialized domains, allowing for the customization of environments and adaptation to human behavior on a much broader scale. In workplaces, these systems will learn the intricacies of professional workflows. Imagine an architect wearing Spatial Computing glasses that not only recognize the materials being used but also predict structural challenges or suggest design improvements based on similar projects from around the world. As the system interacts with thousands of professionals globally, it learns exponentially, improving the workflow for all users in real time.

In medicine, these glasses could recognize symptoms or potential health concerns during everyday activities. A person's biometric data might signal a risk for an illness, prompting the system to suggest scheduling a doctor's appointment or, in more serious cases, contacting emergency services. Imagine healthcare that not only treats illness but predicts and prevents it—a capability fueled by exponential learning and real-time data processing through Spatial Computing.

At the societal level, exponential learning systems won't be limited to personal devices. Entire cities will leverage this technology to become smart cities, where urban infrastructure is continuously learning and optimizing for the needs of its residents. Just as Tesla learns from every braking event to improve safety, citywide networks of cameras, sensors, and AI algorithms will help city planners and engineers make real-time adjustments to traffic flows, energy consumption, and public services.

Imagine a city that learns how to better allocate resources—like electricity and water—based on real-time consumption patterns. If one part of the city needs more energy during peak hours while another is experiencing downtime, the system can balance the load efficiently, reducing waste and increasing sustainability. Public transportation systems will also learn from real-time data to better predict demand, by adjusting bus and train schedules dynamically to meet the needs of residents, reducing overcrowding, and optimizing routes.

These smart cities will function as interconnected systems, where every small event—whether it's a traffic jam, a water leak, or a pedestrian crossing—will feed back into an exponential learning system that continuously adapts to improve the overall quality of life.

With Spatial Computing, AI, and exponential learning systems, we'll also be able to tackle environmental challenges more effectively. Imagine your Spatial Computing glasses detecting waste patterns in your home, suggesting more sustainable practices, or automatically adjusting household settings to minimize energy use. Beyond personal consumption, these systems could identify larger patterns at the community or even global level, providing critical insights to environmentalists and policymakers.

Imagine a future where deforestation, wildlife conservation, and resource management are all optimized through data collection and prediction systems embedded within Spatial Computing networks. By learning from ecosystems around the globe, exponential learning systems will help create solutions that are more targeted and effective, addressing environmental concerns at both local and international levels.

In the coming years, exponential learning systems connected through Spatial Computing glasses and other wearable devices will bring about a new era of connectivity, where personal, professional, and societal growth happens faster and more effectively than ever before. From saving lives with real-time health monitoring to enhancing daily productivity and improving environmental sustainability, the potential of these systems is virtually limitless.

These systems will grow alongside us, getting smarter with every interaction, and over time, they will become integral to how we interact with the world—offering not just convenience but also the ability to reshape entire industries, communities, and ecosystems.

As we continue to connect our lives, our cities, and our planet through these exponential learning systems, we will see the promise of Spatial Computing evolve into a dynamic force that continuously adapts to human needs, making the future more responsive, more intuitive, and more connected than we could have ever imagined.

Data Dance with Seven Verticals

We'll present here some visions of how the intersection of data collection with machine learning and 3D visualization will benefit the seven verticals that we have been addressing in this book: **Transportation; Technology, Media, and Telecommunications (TMT); Manufacturing; Retail; Healthcare; Finance; and Education**. Let's begin by getting into how the world of transportation is set to change due to the influx of data, and our increasing capacity to utilize and learn from it.

Transportation

When you need transportation, you'll see that tons will have changed by 2030. Your Spatial Computing glasses will radically change the way you approach transport. These changes will be bigger than the ones that Uber, Lyft, and Didi brought as they enabled people to order a ride on their new mobile phones. Uber is worth looking back at because we believe transportation, even the vehicle that sits on your own driveway, will act and work a lot more like Uber than like the Chevy that was memorialized in the 1971 anthem, "American Pie."

When Travis Kalanick was starting Uber, he told us that his goal was to have a single car near you. Not none, because that would mean waiting, which wouldn't be a good customer experience. Not more than one, because that would mean drivers were sitting around not getting paid and getting angry.

This is accomplished by data collection, smart data analysis, and proprietary algorithmic software. It knows, contextually, where rides are. In fact, drivers are urged to move around the city, and sometimes further, in anticipation of new demand. We watched its organization at the Coachella music festival, where drivers from across the Western United States—thousands of them—converged in Palm Springs by this system that controlled transportation.

The transportation system of 2030 will be different, though, because most of the pioneers we talked with expect autonomous systems to be well underway. In fact, Waymo is driving people around in these, now named "robotaxis." At least it is in a few cities, like Phoenix, Arizona, and Mountain View, California. We cover these robotaxis in depth in *Chapter 4, Vision One – Transportation Automates*, but it's worth explaining how the data that these systems are slurping up everywhere will be used, processed, and displayed, all in an attempt to predict what the transportation needs are so that Kalanick's impulse of having a vehicle waiting nearby whenever you need it will become true, even without a driver involved.

Further, your glasses, and possibly Augmented Reality windshields, will see layers of other utilities, all built by the data being gathered. For instance, it's very possible that you will be able to see what the street looks like during the day from these glasses or windshields. The maps that are being built by autonomous cars are so detailed and have so much data about what surrounds the car that you might be presented with a "day mode" so that you can more easily see what surrounds the car.

The maps themselves, which the industry calls "HD maps" (for "high definition"), already go way beyond the maps most of us use in mobile apps, whether from Google or Apple. They include every stoplight, stop sign, and tree on the route, lane information, speed information, and much, much more.

The autonomous car system needs all that data to properly be able to drive you from your home to, say, work—stopping correctly at intersections, accelerating to the right speed on freeways, and turning smoothly on mountain roads. All this is controlled by the HD maps underneath, usually invisible to the driver. However, if you are in a recent Tesla model, you get a taste of these maps on the car’s display, which, as we are writing, now shows stoplights and even things on the road like garbage cans. The data these maps are collecting, though—or rather, that the autonomous car groups, whether Cruise at General Motors, Waymo, or Tesla, are gathering—goes way beyond what you might think of as data for moving vehicles around.

Because they need to track literally everything on the road, their AI systems have been taught to recognize everything from a ball to a child riding a bike, and their systems predict the next likely action of each of these things. These systems have also learned how to recognize parking meters, parking spaces, curbs, poles, and other things that they might need to navigate into and out of.

Now, think about when Spatial Computing glasses and autonomous cars work together. The car system itself can watch how many people are waiting in line at, say, a nightclub. It can tell how busy a downtown district is, or whether a restaurant is either quieter or busier than usual.

At Microsoft, 15 years ago, they showed us how predictive systems could use huge amounts of data to predict just how busy a freeway would be at a certain time, long before you would leave on your journey. Today, that kind of prediction seems pretty quaint as Spatial Computing glasses and autonomous vehicle systems are gathering 10,000 times as much, or more, data than was gathered by early AI-based systems, and the amount of data is increasing exponentially every year.

So, what is possible as these systems mature and as more and more data is gathered by fleets that grow from hundreds of thousands of AI-outfitted fleets to millions of cars, all assisted like an ant farm with millions of people gathering other data on their wearable computers that are also mapping the world out in 3D?

You’ll be able to see literally everything about the street in real time. Is there a wreck ahead? You’ll see it within seconds of it happening, and potentially even be able to watch it happen in 3D on your glasses. Lots of other news will also be displayed on maps in the future in this way—everything from bank robberies to fires. Already, lots of cars are able to display video captured from dashcams or the cameras included in, say, a Tesla, which has AI that senses that humans are moving around the car in a way that they shouldn’t, so it starts playing loud music and records them in an attempt to get them to move away from the car.

Systems like Waze have long warned of things like potholes or objects on the road ahead, but in a world where cars have 5G radios and can transmit huge amounts of 3D data, by 2030, you will be able to see changing road conditions sent to you by cars imaging those conditions ahead. The thing is, most of us around that time won't care about road conditions at all, except maybe in some extreme weather situations where sensors stop working, or the road can't properly be sensed. Even then, the autonomous car network might "hand over" control of your vehicle to an employee wearing an expensive Virtual Reality headset and they could remotely control your car without the passenger even knowing about it—all thanks to the data systems that are gathering, processing, and sharing a 3D view of both the real world and a digital twin, or copy, of the real world.

Speaking of that digital twin, with your Spatial Computing glasses, you will be able to see a whole raft of new kinds of games, along with new kinds of utilities. Pass by a historical marker? Well, grab a virtual version of that with your hand and pull it closer to you so you can read the sign, all while your car zooms past. Want a hamburger? Tell your glasses your needs and you will be satisfied not just with a "McDonald's is 2.5 miles away" message but with a full 3D menu of the food options so you can have your order waiting for you when your car navigates there and you'll know exactly what to expect, whether it's a salad, a hamburger, or something more complex like a Vietnamese soup.

The kids, too, will have their own Spatial Computing glasses and will invite you to join in their game, where they are manipulating a metaverse, or virtual world, that you are driving through. Toyota, and others, have already demonstrated such Virtual and Augmented Reality games that are under development.

It is the predictions that are possible because of all the data being collected as you navigate, stop for food, go to church, and head to work that are really set to shake things up. Remember, a system that has hundreds of thousands of cars, each with eight cameras and a variety of other sensors, can "see" patterns humans just can't. The system can predict what's next, knowing that a baseball game is about to finish, for instance, and can reroute everyone within seconds. Humans just aren't that good at knowing things like that. These predictions, done on a mass scale, will save millions of people a few seconds a day, sometimes minutes. Add those all up and it's the same as saving a life's worth of time every day, and at a bigger scale, it will truly revolutionize transportation.

These predictive databases and the AIs that are working on your behalf will make the transportation system itself much more efficient, safer, and more affordable. For instance, if you were to take a car on a long trip, these systems would know that long before you left, probably because you put something about the trip on your calendar, or told social media friends "Hey, we are going to New Orleans next weekend.

Anyone want to meet up?" Or, even without any signals, it would figure out that you were on a long-range trip as soon as you crossed a geofence 30 miles out of town. All the way, the master in the sky would be able to route you the fastest way and even hook you up in a convoy to save everyone some money on electricity for your vehicles.

Fleets like those owned by UPS earlier started to move toward being completely automated and electric. In 2020, UPS announced that it would buy 10,000 electric vans and install Waymo's autonomous driving package onto each (Waymo is the autonomous vehicle start-up spun out of Google's R&D labs, and was the first self-driving vehicle system we saw on Silicon Valley roads). Waymo paused this project in 2023 to focus on its robotaxi operations.

The predictive system taking care of you on trips, and preparing the car for you automatically, saves a few seconds of typing an address into a phone-based map or clicking around on a screen to get it to navigate, an act that can be very dangerous if you attempt it while driving, which many people do. It seems like something minor, but over the next decade, our glasses and car screens will add dozens, if not hundreds, of little features like these to make our lives easier.

After all, we are creatures of habit. Sam Liang, who used to work on the Google Maps team, told us just that. He recognized years ago that the context in which humans were trying to do something matters to make computing systems work better.

After all, if you are wearing your glasses while driving a car, you will expect a different set of functions to be available to you than, say, when you are attending church, and that will be different than if you are watching a movie.

The cloud computing infrastructure that is watching your glasses, and all the cars moving around you, is about to make life itself much different than it was a decade ago when Uber seemed like a big idea. Let's discover how it'll work, not just when you are ordering a car but in the technology you utilize, the media you consume, and the communication systems you use.

Technology, Media, and Telecommunications (TMT)

Where and what you pay attention to is already paid a lot of attention. Teams of engineers at Netflix, YouTube, Hulu, Spotify, Disney, and other media and entertainment services have developed systems to pay attention to what media choices you make and where and what you watch, and they have already built extensive prediction systems. You can see these systems come to life when your kids use your accounts for the first time and all of a sudden you start seeing suggestions that you watch more SpongeBob.

In the Spatial Computing world, though, the data collected will go far beyond what movie, TV show, or website link you click on. There are seven cameras on the devices that Qualcomm's XR2 chipset enables: four for looking at the world, one for watching your face for sentiment, and two for watching your eyes.

It is the ability to see user sentiment and watch where users are looking that will open up many new predictive capabilities that will change even the entertainment we see. Knowing you actually looked at a door in a game or interactive movie of the future, for instance, could put you down a different path than if you looked elsewhere. Thanks to the incredible download speeds that wireless will soon bring with 5G, the device you are wearing itself could load quite complex 3D models before you even open the door to look through.

Games like Simulator 2024 are loading assets based on where you are located, and if you walk down the street, it predicts you will continue walking down the street and loads characters and game pieces ahead of when you need to see them.

Microsoft's Flight Simulator has worked this way for more than a decade. As you fly your virtual plane, the system loads scenery and city details ahead of you so that everything is sharp and doesn't stutter when you fly over, say, New York. Entertainment makers are planning out new kinds of entertainment that let you choose a variety of paths through the entertainment. Netflix already has shown how this will work with its TV project "Black Mirror: Bandersnatch." There, you can make choices for the characters that shape the story as you go. In Spatial Computing glasses, touching things, looking at other things, and moving around your real-world living room might impel movies and games you are enjoying to show you different story paths.

Lots of pioneers in the industry refer to an off-Broadway New York play, "Sleep No More," which was a new version of Shakespeare's Macbeth, except it wasn't performed on a stage. Audience members wandered through a huge warehouse, with action happening all around as they walked through the set.

Innovators like Edward Saatchi, who is the co-founder and CEO of Fable Simulation, see this as foundational to the future of entertainment. Soon, he says, audiences will be walking through virtual sets, similar to in "Sleep No More," and will be interacting with virtual beings that can talk with you, or even play games with you, among other interactions. Fable's "Wolves in the Walls" is an example of just this and introduces a virtual character, Lucy, who can do just this. This Emmy-winning project has Lucy talking and interacting with you and is the first example of a virtual being we've seen that makes you feel like you are inside a movie.

These virtual beings will be programmed to interact with you, asking you questions, presenting you with choices, and constantly feeding an exponential learning machine running behind the scenes. “Would you like to listen to Taylor Swift or Elton John?” she might ask one day. The answers you give will change every interaction you have from then on, and you might not even know, or care, just how deeply your experience has been manipulated, or how much a company behind the scenes is learning about your preferences and passions.

Give an exponential learning machine enough of these interactions and it quite possibly could figure out that you are a conservative Republican, or that you will soon be a full-blown country music fan, or even that you are of a particular sexual orientation—even if you’ve not yet been entirely honest with yourself about it. Each can be used to get you addicted to more things that interest you, but there’s a dystopian fear that the system could, even by accident, divide you from other humans and radicalize you into having potentially negative beliefs. We’ll discuss that more in *Chapter 15, How Human?*

For now, know that Spatial Computing glasses are arriving at the same time as extremely high-speed wireless networking and AI that can drive cars and predict the future moves of hundreds of people and animals around you. Add this all together and you’ll be experiencing new kinds of movies and games and will see the real world in quite different ways with many augmented layers that could show various things both above and below the street, including other people’s information.

Don’t believe us? Meta has a strategy book that it built for the next five to ten years, where it details how people will use Spatial Computing and describes how you will know whether someone is a friend, or connected to a friend, on the street, and enable you to play new kinds of video games with them. Already in Virtual Reality in the Meta Quest 3 headset, you can play paintball and basketball, among dozens of other games, with friends and random people. This can be done using Meta’s Horizon Worlds app, as well as playing any of the thousands of independently made games.

Next, let’s think about how Spatial Computing and the data associated with it will impact the world of manufacturing.

Manufacturing

We predict that within the next five to seven years, Virtual Reality and Augmented Reality combined with computer vision and machine learning-massaged data will be used much more in manufacturing than they are currently. The drivers for making this happen have to do with efficiency, quality, and safety.

Efficiency, quality, and safety will all increase when more Spatial Computing systems are integrated into manufacturing processes because the humans working in manufacturing could then focus less on the repetitive aspects of their jobs that are currently taking up much of their time, and focus on other tasks where human decision-making processes could be used. The negative aspects of human rationale when encountering repetitive tasks could be avoided—those associated with a lack of focus and tiredness that lead to slower productivity, lower quality assurance, and injuries.

Additionally, these Spatial Computing systems, most notably those using Augmented Reality, will be fully integrated to work with robotics and other machinery so that currently existing setups can be improved upon. Besides being used for making products, both Virtual Reality and Augmented Reality will be used for training purposes, as well as for accessing machine breakdowns and recommending solutions.

Porsche Cars North America has rolled out an AR-based solution for service technicians at dealerships, for instance. Why? Porsche and its vendor that built the system, Atheer, say that the system connects dealership technicians to remote experts via smartglasses for a live interaction that can shorten service resolution times by up to 40 percent.

In complex assembly, using AR sees even bigger gains. As Boeing builds a 747-8 Freighter, its workers keep track of 130 miles worth of wiring with smart glasses and the Skylight platform from Upskill. The result? Boeing cut wiring production time by 25 percent and reduced error rates effectively to zero.

At BAE, they are using Microsoft HoloLens headsets to build the motors for their electric buses. The software system they used, Vuforia Studio from PTC, let them train new people 30 to 40 percent more efficiently and build these motors much faster and with lower error rates.

Once a few more large corporations adopt Spatial Computing in their manufacturing operations, they will become even more operationally efficient, which will lower their operational costs and give them market advantages. This will cause other companies to adopt Spatial Computing technology for their manufacturing operations until it becomes a business standard.

Having covered manufacturing, let's now move on to think about the world of retail.

Retail

As you walk into Sephora's R&D lab in San Francisco, you see the depth to which they are thinking of using technology to improve the customer's experience there. They built an entire copy of the store so they could run tons of customers through the store and study every interaction.

This shows up on perfume displays that have been digitized, so you can smell each product without spraying it on your skin, along with AR-based signs, and even virtual makeup that you can try. All of these experiences are improved through the use of data on all sorts of things that happen in the store, from the time that customers spend hanging out in front of displays to real-time sales that provide feedback loops, to promotions and store changes.

Today's e-commerce stores have trouble vividly showing you how a product will actually look when you receive it. Spatial Computing glasses change all that by bringing you 3D visuals of products, and we predict 3D stores where you will be able to grab products from virtual shelves, spin them around in your hands, and even try them out.

This will give retailers whole new sets of data to study and use to improve their presentations and the products themselves, but it is the predictive possibilities of these new displays that have the retail industry most excited, said Raffaella Camera, former Global Head of Innovation & Strategy, XR and Managing Director at Accenture Interactive.

What do we mean? Well, let's say you add a wedding or a baby shower to your calendar; these systems can bring new shopping experiences to life. Imagine a wedding registry populated with 3D products that your friend has picked out, along with maybe 3D images of your friend explaining why they added that item. Touch an item or show some interest in it, and it can change the entire display. For instance, touch a china plate and it could instantly show you all sorts of other things that pair with it, from glassware to silverware that you might also purchase for your friend.

The industry is salivating at getting this new data, which includes where you look, what you touch, and what you buy.

“Give me eye-tracking data of you in the real world for a week, just going about your business,” said Roope Rainisto, former Chief Design Officer and co-founder of Varjo, which makes the high-end VR headset that Volvo is using to design new car interiors. “The value of that data for learning about your preferences is so immense—behaviors that are currently not visible to these online services.”

What will that data be used for? Rainisto sees a new shopping assistant. “There is far too much information for one person to understand and to even keep track of. The personal assistant will swoop in, help us by telling us things we should know or would like to know, do menial tasks, like keeping track of the prices of particular goods, put items on a wish list for the future, etc.”

“How does the assistant become good at its job?” he asks. “Through learning us. This learning is fundamentally the exact same learning as what is used in ad targeting. Once the assistant knows me, then it can properly help me. Before that, its services will ‘feel like ads’—pushing stuff I do not find all that relevant or helpful.”

Retail itself will be changed deeply by these predictive systems and by workers who use Spatial Computing devices. For instance, today, if you buy something in, say, B&H Photo in New York, an order goes into a warehouse and a worker grabs the camera you want and puts it into a tray that goes on a conveyor belt to where you complete your purchase. If you are buying online, the same process happens, except your product goes to a separate packing warehouse. Soon, these systems will be made more efficient thanks to workers who are directed on what to do by wearable computers that can see the products via computer vision and show the worker all sorts of details via AR.

These warehouse workers using Spatial Computing devices and hooked into a system of retail helpers (both real humans and virtual assistants) will radically change retail.

Logistic robots are already doing this, but not as efficiently as humans because computer vision for robots, along with ML, is not sophisticated enough to recognize novel objects and objects that are not uniform. There are researchers already working to make this possible, but it will take another five years.

When this happens, very few humans will be working in logistics. Even with today's limited technology, this process of automating everything is underway; we've met warehouse owners who tell us they turned on completely automated warehouses as early as 2019.

Like with transportation, retail too will have a cloud computing system watching literally everything. Soon, such a system might even save your life; let's talk healthcare.

Healthcare

The combination of data, ML, and 3D visualization will be a great boon to healthcare. The data and ML combination have been used by practitioners, as well as non-practitioners, as in the case of the Apple Watch, for a few years already. The inclusion of 3D visualization, paired with location and situational awareness, makes a big difference, even saving a few lives as the systems watching sensors can tell whether you have taken a terrible fall or are having a heart attack.

Healthcare practitioners will be able to view patients' current conditions, including 3D scans of their internal organs, as well as run simulations of conditions to plan and predict how surgery and other kinds of interventional procedures should go for those patients. As we mentioned in *Chapter 8, Vision Five – Virtual Healthcare* on healthcare, a company called MediView is already using AR actively during cancer ablation surgeries. In the future, decisions about whether or not to perform surgery could be made using Spatial Computing methods.

The big change that will enable this to happen is the massive digitization of patient data that is currently underway now by hospitals and other organizations. Issues regarding this are the different systems that are being used by the different organizations, how they might not be universally accessible, and privacy issues having to do with bulk data. Without data being generally accessible, applying ML will not provide trustable outcomes. We believe that, eventually, perhaps by the mid-2030s, medical data sharing via patient opt-ins will become more ubiquitous due to the immense benefit to society that would occur as a result.

Moving on from healthcare, let's take a look at how change is coming to finance.

Finance

We envision a day when many types of financial analyses and trading will be accomplished using Spatial Computing glasses. Much of this could be accomplished because of the massive amounts of raw data streams that will be made available to make decisions, as well as data that has been characterized by ML. All of this relevant data could be portrayed in 3D so that decision-making could be made even faster.

Virtual assistants could also be used in combination with voice querying, as well as proactive and opted-in suggestions honed in by ML to aid in making decisions as to whether or not to participate in a trade.

In addition to financial analyses and trading and 3D data representations, Spatial Computing can also be used for customer-financial service rep interfaces. The combination of the use of voice, data, ML, and 3D visuals could make figuring out thorny financial issues much easier and faster.

Add in cameras and computer vision, and now we have facial recognition technologies that use ML, as well as body-gait analysis technologies. With these, security is greatly increased, and all kinds of financial transactions could be more efficiently accomplished.

Finance is truly one of those areas that will be completely disrupted by the use of Spatial Computing; as it is with something like this, it is difficult to completely foresee all the different kinds of new apps and technologies that will emerge as a result of Spatial Computing's disruptive force.

Education

Millions of people remained unemployed for a significant amount of time post-COVID-19, and the coming increased automation in manufacturing and cars will bring on even more upheaval in society. Truck drivers, for instance, will largely lose their jobs in the next few decades. To America alone, that means a pretty deep restructuring of the landscape because trucks will no longer need to stop at small towns along interstates, which allow hotel keepers and cafe workers alike to be employed.

Millions of people will see their livelihoods threatened and will need retraining to remain relevant in the modern world.

Even if you aren't struggling to put food on the table, though, Spatial Computing will radically change what it means to "be educated." Everyday activities like travel, cooking, learning about the latest science, watching the news, and even participating in civic duties will change deeply because of Spatial Computing. We can see education changing dramatically in our discussions with founders of companies who are already planning ahead with 3D lesson plans and are building exponential learning engines to teach us everything from how to save a life to how to cook a certain meal. Our original June Oven already had a camera and an NVIDIA card, running AI so it could automatically detect the kind of food we placed in it and properly set the cooking temperature and time. On its app, it was already exposing us to an exponential learning engine that got better over time, showing us more and more recipes, and properly recognizing more and more food types, from different kinds of fish and meat to different kinds of bread.

In *Chapter 10, Vision Seven – Real-Time Learning* on education, we cover how extensively the field is being changed by Spatial Computing. Here, we explore some things that are coming due to always-connected exponential learning engines, which will soon change every part of our world, including how we learn.

For instance, in China, students are being checked in via facial recognition. In other places in China, teachers are starting to be replaced by AI and VR. Well, actually, not replaced, since there aren't enough teachers to keep up with demand. Soon, the same might happen in the USA, and what we call exponential learning systems, which are a combination of data and applied AI that constantly learns from an increasing collection of data, will take over the act of educating.

Why? We won't have enough teachers to keep up with demand soon, either. Look at truck driving, which is the number one job in the United States with 1.3 million people driving today. What will happen over the next decade or two when their jobs go away? Who will teach them to do something else? Already, companies like Caterpillar are using AR glasses to teach people how to fix tractors in real time through visual overlays on top of the tractor itself. Mercedes-Benz is using AR to teach first responders how to cut apart a car without touching an electric line or fuel line. This app, designed by Germany's Re'flekt AR development firm, works on a tablet today and glasses tomorrow and shows firefighters arriving on the scene at a car wreck where it's safe to cut.

The use of AR alone will make education better, but think about any task. Could it be made better by hooking it up to an exponential learning machine? Even the Mercedes app could be made better, by studying, say, how a firefighter in Japan uses it on a custom rebuilt car. That learning could be shared with the cloud computer in the sky and shared instantly with all the other firefighters in the world, making everyone's skills and knowledge better.

Can such a system learn via ML that humans aren't doing a task right and gently nudge them with visuals in their glasses that there is a better way to do something? Even something as difficult as, say, learning a language, how to play an instrument, or how to make a complicated meal can be made better by joining an exponential learning engine with Spatial Computing glasses displaying VR and AR to its wearer, and mundane things like cooking an egg or replacing a tire are made much simpler with AR assistants. Yes, you can look those up on YouTube already, but if you were wearing Spatial Computing glasses, you would be able to see assistant images and directions right on top of what you are trying to learn, and the sensors on the glasses could watch to make sure you took the right steps, too.

Classroom teachers are currently getting a taste of Augmented Reality with the Merge Cube, a low-cost plastic cube you hold in front of a camera on a phone or tablet that overlays AR images. But this is a very elementary example of what could be accomplished with Spatial Computing in the future. Keeping you immersed in the digital world makes learning better, and letting you use your hands on top of a real task will enable new education that is hard to imagine today. Like the employees at Caterpillar, students will learn through information and media superimposed on top of the task they are trying to master in real time.

Imagine learning chemistry or surgery this way! It's already happening, but this will speed up dramatically over the coming decade and affect every part of human life, always making sure we have the latest knowledge and visualization to help us understand in 3D, which is how our brains think and learn—all with data that's constantly getting better and more detailed. Already, AI is finding cancer in medical scans more accurately than human doctors can. In the future, an exponential learning engine will connect all scanners and doctors, and all doctors utilizing the system will receive the benefits each time the system learns something new.

Can schoolchildren be taught by such a method? While we can't see a world where an adult isn't nearby or facilitating learning, we can see many lessons taught via AR and VR. When we were in school in the 1970s and 1980s, many lessons were taught via film strips or movies. Soon, those kinds of lessons will be taught by asking children to put on a pair of "learning glasses." These glasses will teach everything from history to chemical reactions, and the connected exponential learning systems will keep track of every lesson and every student. The adults in the room could then still be notified if someone is falling behind, falling asleep, or goofing off, which is something we still see children doing in the future.

We've now considered the seven verticals. Let's zoom out to a wider perspective in the next and final section and think about something that we refer to as the "bubbleland."

The Bubbleland

We soon will live in a digital ant farm, or, as we put it, the “bubbleland.” We came up with the term after watching how autonomous cars see the world. Each has a digital bubble around it, beyond which it cannot see. Humans work the same way, but we don’t think of our perception systems that way, sometimes believing that we can even sense the unseen. In the digital world of Spatial Computing, in the case of glasses, robots, virtual beings, and autonomous vehicles, however, their perception systems are limited to what their cameras and 3D sensors can see.

Are they really limited that way? If you throw millions of such little bubbles onto a city, are there many places that a bubble doesn’t touch? No.

The Bubble Network

We see such a city as having an exponentially growing number of exponential learning machines, each with a bubble of what they can sense, all communicating with a central cloud computing system. Add a bubble or a new set of glasses or a new autonomous vehicle and this unified system gets more information that it can use to “see” or understand the world we live in better. Add even more and we’ll get more services and more utility. It is a new network effect. In previous generations of networks, as more people joined in, the network got more interesting. This is why the web eventually got so complex we needed search engines like Yahoo and Google to keep track of it all.

The new network effect, though, isn’t about the information on pages; it’s these information bubbles moving around the real world. One such bubble might have to do with all the information circulating while waiting in line at a nightclub. Another might be around someone pumping gas into their car. Another might be around those worshipping at church. Yet another could occur when picking up groceries for tonight’s family dinner.

Unified Learning, Shared Experiences

A unified exponential learning system can keep track of all of this, seeing inventory changes at grocery stores, lines at gas stations, activity at nightclubs, keeping track of the latest learning delivered at universities, and much more. Each bubble would be training the system and teaching AIs what to further look for. This system eventually will turn even hungrier for data and might ask the people wearing the glasses or the robot car to move around and gather more data. “Hey, we haven’t seen aisle three in a few hours, can you walk down there with your glasses and check out the cookie display?” our glasses might ask soon, maybe even offering a discount on a purchase to encourage us to provide more data so that the system has accurate cookie prices and inventory data to share with everyone on the network.

Soon, maps will show these bubbles in real time, or at least the data that they are gathering. Today, when you see an accident on the road, it shows a red line to denote that traffic is backed up, or an icon where the accident is reported. Tomorrow, you will see a 3D version of the wreck on the map and you will know the speed of each lane of traffic as reported by these bubbles.

It'll be the same for grocery stores, churches, malls, movie theaters, restaurants, laundromats, national parks, and more. You might even be able to "dive into" a bubble that is explicitly sharing its digital bubble with the world. Wouldn't you love to do that in someone's bubble who is in the front row of a Lady Gaga concert? We know we would, and we would be happy to share such an experience by telling our glasses "Hey Siri, share this with the world." Other experiences might be more specific: "Hey Siri, share this with Grandma" or "Hey Siri, share this with other parents at our school."

The Spatial Computing glasses and the autonomous cars and robots of the future will take the concept a lot further, to where they teach exponential learning systems a lot more about the world as they are walked around, and present viewers with tons of interesting views of the real world, including new digital twins that someone thousands of miles away will be able to move around in, interact with, and maybe even add to. This bubbleland of the future is one that is very human, creative, and inevitable.

Pervasive and Proactive Future Assistance

As we've shown here, soon, our technology will know we haven't yet put out the garbage or will prepare you for your next meeting in ways you never thought of, not to mention perhaps warning you about a health problem that's coming for you based on your biometric data. Our Tesla cars sense when garbage cans are out, so this day isn't far away, and soon, not only will Teslas perform this kind of object and pattern recognition but so will millions of people wearing new Spatial Computing glasses, and little delivery robots rolling around stores and streets.

Next, let's look at a few of the pioneers who are already starting to use these technologies to change everything from transportation to shopping.

12

The Business of Generative AI

The fusion of Generative AI and Spatial Computing is reshaping the business landscape, offering new dimensions of interaction, creativity, and operational efficiency. Companies are leveraging these technologies to develop innovative products, enhance customer experiences, and streamline processes, positioning themselves ahead of competitors in an increasingly technology-driven marketplace. This chapter explores the applications, economic impact, challenges, and future potential of these technologies within various business sectors.

Generative AI and Spatial Computing are two transformative technologies that, when combined, provide unprecedented opportunities for businesses. Generative AI, a subset of AI that focuses on creating new content—such as images, text, music, or even entire virtual environments—opens doors for innovation in content generation, product development, and automation. Spatial Computing, on the other hand, integrates physical and digital worlds through technologies like AR, VR, and 3D modeling, creating immersive experiences and enhancing human-computer interactions.

The intersection of these technologies enables businesses to design unique digital experiences, prototype products faster, and create personalized marketing strategies. Understanding the business implications of this convergence is essential for companies aiming to capitalize on the advantages these technologies offer.

Generative AI in General

Generative AI, a branch of AI, is transforming industries by harnessing the power of deep learning models to create new content based on patterns from existing data. This technology leverages neural networks, particularly **generative adversarial networks (GANs)** and transformers, to generate outputs that can be images, text, music, 3D models, and even entire virtual environments. The applications of Generative AI are wide-ranging, influencing fields from content creation and product development to customer service and creative industries. This section explores these applications in detail, demonstrating how Generative AI is reshaping business operations, enhancing efficiency, and enabling new revenue streams.

Content Creation and Media Automation

Generative AI is revolutionizing the way content is created, optimized, and distributed across various platforms. In the marketing and advertising sector, AI models analyze vast datasets that include user behavior, preferences, and demographic information to generate highly targeted and personalized advertising copy. Instead of creating generic ads, businesses can now produce copy that resonates with specific audiences, enhancing the effectiveness of marketing campaigns and increasing conversion rates. Furthermore, AI-generated visuals for social media, including custom graphics and video snippets, enable companies to maintain a consistent online presence while reducing the need for human designers.

News organizations and media companies are also increasingly relying on AI to automate the creation of written content. AI systems are capable of generating articles, news summaries, and even full-length reports by processing large volumes of information in real time. For example, sports reporting and financial updates, which often follow structured formats, are now frequently produced by AI, ensuring timely and accurate delivery. This automation not only reduces the workload for journalists but also allows media outlets to scale their production capabilities, covering more stories with fewer resources. In addition, AI can generate graphics and infographics that accompany articles, further enhancing the visual appeal and informativeness of content.

Generative AI is also being used in the field of entertainment and streaming media. Personalized playlists, video thumbnails, and even scriptwriting for short videos are now driven by AI models that understand audience preferences and engagement patterns. Streaming platforms leverage these capabilities to suggest highly relevant content, thereby increasing viewer retention and satisfaction. For instance, AI-generated trailers can automatically highlight key moments in films or shows, tailored specifically to individual viewers' tastes, ensuring higher click-through and engagement rates.

Product Design and Prototyping

Generative AI is reshaping the product development process across industries such as automotive, fashion, and architecture. Traditional product design often involves time-consuming manual iterations to refine concepts and improve efficiency. However, AI-driven design software can analyze performance criteria and aesthetic requirements, automatically generating multiple design variations based on specific parameters that optimize these factors. In the automotive industry, for instance, engineers can use AI to generate car models that not only meet aesthetic preferences but also incorporate aerodynamic features and material efficiency, reducing the time required to move from concept to prototype.

Fashion companies are increasingly turning to AI to accelerate their design cycles. AI tools can generate clothing patterns based on fashion trends and historical data, offering designers a range of new styles to choose from or refine further. This technology not only shortens design timelines but also reduces material waste by suggesting optimal fabric usage patterns. Fashion brands are also using AI to create virtual try-on experiences, where digital models of new clothing lines are tested virtually before physical production, further streamlining the prototyping process.

In architecture, Generative AI assists architects by automating the creation of building models that balance aesthetics, functionality, and environmental sustainability. AI-driven design platforms can analyze structural requirements, climate conditions, and urban planning regulations to propose building designs that optimize for these parameters. This approach allows architects to experiment with innovative designs that would be difficult or time-consuming to develop manually. Moreover, these AI-generated models can be visualized in AR/VR environments, enabling architects and clients to explore and refine projects in immersive 3D spaces before construction begins.

AI-Driven Customer Interactions

Generative AI is also transforming customer interactions through the use of AI-powered chatbots and virtual assistants. These systems are now capable of providing personalized and dynamic support, responding to queries, and offering recommendations based on user data and historical interactions. Unlike traditional scripted chatbots, modern AI-driven systems utilize **natural language processing (NLP)** and **machine learning (ML)** algorithms to understand context and engage in more human-like conversations. This capability not only enhances user experience but also reduces operational costs, as companies can manage large volumes of customer interactions without the need for human agents.

For businesses operating in global markets, AI-driven virtual assistants offer multilingual support, enabling companies to provide consistent customer service across different regions and languages. This flexibility is crucial for global retail, hospitality, and tech companies that need to manage inquiries from diverse customer bases. Additionally, AI systems can learn and adapt over time, improving their responses and refining their customer engagement strategies. For example, if a user frequently asks about a specific product category, the AI system can proactively recommend new products or promotions related to that category, increasing the likelihood of a sale.

Generative AI also plays a significant role in personalizing customer interactions through recommendation systems. E-commerce platforms and streaming services use AI algorithms to analyze user behavior and preferences, generating personalized product or content recommendations. These AI systems enhance user engagement by suggesting items or media that are highly relevant to individual customers, improving both the user experience and sales performance. In the hospitality industry, AI-driven concierge services offer tailored recommendations for activities, dining, and services based on guests' past preferences and real-time data, such as weather conditions or local events.

Expanding Creativity in Entertainment and the Arts

Generative AI is having a transformative impact on the creative industries, including music, art, film, and literature. In music production, AI algorithms are capable of composing original tracks or collaborating with human musicians to enhance compositions. These AI-generated tracks can be tailored to fit specific moods or styles, making them valuable for advertisements, film scores, and video game soundtracks. Companies like Suno and Brev.ai are leading the way in developing AI systems that assist composers and music producers, expanding the range of creative possibilities while reducing production time.

In visual art, AI models like DALL-E, Midjourney, Leonardo, and Runway use deep learning to generate images and/or video and artwork based on text inputs or by blending elements from different styles. These tools enable artists to explore new creative directions and experiment with styles that may not be possible through traditional methods alone. For commercial purposes, businesses are using AI to generate marketing visuals, design patterns for fashion, and create digital art for video games or animations. AI-generated art also opens up new revenue streams, as companies can sell these pieces as **NFTs (non-fungible tokens)**, tapping into the market for digital art ownership.

In the film and gaming industries, Generative AI is pushing the boundaries of storytelling and visual effects. AI-generated scripts, dialogue, and even entire scenes are becoming more common, enabling filmmakers and game developers to experiment with narrative variations quickly. AI models can generate realistic 3D characters and environments, allowing for the rapid prototyping of scenes and the creation of immersive experiences that would otherwise require extensive manual labor. For example, AI can create realistic facial expressions and body movements for digital characters, enhancing the realism of virtual environments in VR games and films.

Literature is another domain where Generative AI is making strides. AI models like **GPT (Generative Pre-trained Transformer)** can generate entire stories or assist writers by suggesting plot developments and character interactions. While AI is not yet capable of replacing human authors, it serves as a valuable tool for brainstorming and enhancing storytelling. Publishers are also using AI to generate personalized content for readers, such as customized story endings or interactive books that adapt based on reader choices, providing a unique and engaging experience.

Revolutionizing Education and Training

Generative AI's influence extends beyond creative industries into education and training, where it is reshaping the learning experience. AI-driven educational platforms are now using generative models to create interactive content, such as personalized lesson plans, quizzes, and virtual simulations. By analyzing student performance and learning preferences, these systems adapt educational materials to suit individual needs, improving engagement and learning outcomes. For instance, AI can generate customized problem sets in mathematics or personalized reading comprehension passages that align with a student's interests and skill level.

In professional training, Generative AI is used to develop virtual simulations and immersive training modules in VR environments. For industries like healthcare, aviation, and manufacturing, where hands-on experience is crucial, AI-generated simulations provide realistic and safe environments for trainees to practice their skills. For example, medical students can use VR simulations powered by AI to perform virtual surgeries, gaining experience without the risks associated with real-life operations. Similarly, pilots and engineers use AI-driven simulations to practice emergency scenarios, ensuring that they are better prepared for real-world challenges.

Corporate training programs also benefit from AI-driven customization. Businesses can create training modules tailored to individual employees' roles, skills, and development goals. Generative AI generates relevant case studies, interactive scenarios, and role-playing exercises that allow employees to gain practical experience in a controlled, immersive environment. By offering adaptive training programs, companies can enhance the effectiveness of their learning initiatives while reducing costs associated with traditional, instructor-led sessions.

Healthcare Innovation and Medical Research

Generative AI is making significant contributions to healthcare by enhancing medical research, diagnostics, and patient care. In drug discovery, AI models analyze vast datasets to identify potential drug candidates, predict their effectiveness, and optimize molecular structures. This approach accelerates the research process, reducing the time it takes to develop new treatments and bringing life-saving drugs to market faster. AI-generated models are also used to simulate drug interactions, helping researchers predict side effects and refine treatments before clinical trials begin.

In medical imaging, Generative AI assists in creating highly detailed scans and enhancing the accuracy of diagnostics. AI algorithms can generate enhanced 3D models from CT scans or MRIs, allowing doctors to explore patient anatomy in greater detail. AI-driven image recognition tools detect anomalies, such as tumors or fractures, with high accuracy, supporting healthcare professionals in making quicker and more accurate diagnoses. The use of generative models in medical imaging also extends to reconstructing incomplete scans, providing a clearer view of medical conditions without the need for additional imaging procedures.

Generative AI is also used to personalize treatment plans. AI systems analyze patient data, including genetic information, lifestyle habits, and medical history, to generate customized treatment strategies. In fields like oncology, AI can suggest personalized chemotherapy protocols based on the specific genetic mutations of a patient's cancer cells. This level of customization improves patient outcomes and reduces the risk of adverse reactions, paving the way for more effective and personalized healthcare solutions.

Generative AI is a transformative technology with broad applications across industries, including content creation, product design, customer interaction, creative arts, education, and healthcare. By leveraging ML and deep learning techniques, businesses are automating processes, enhancing productivity, and generating new sources of revenue. As this technology continues to evolve, its potential to reshape industries and create innovative solutions is likely to expand further, providing businesses with unprecedented opportunities for growth and development.

The Convergence of Generative AI and Spatial Computing

The convergence of Generative AI and Spatial Computing is a revolutionary force in various industries, unlocking transformative opportunities that reshape how businesses operate, create, and engage with customers.

By leveraging these two technologies, companies are not only enhancing efficiency and personalization but are also building entirely new business models and immersive experiences. The integration of these technologies goes beyond mere technical innovation; it changes the dynamics of industries like entertainment, retail, urban development, marketing, and more.

Automated Content Generation for Immersive Environments

The integration of Generative AI with Spatial Computing has opened up transformative possibilities for creating immersive environments, particularly within AR and VR platforms. This fusion of technologies enables the automated generation of complex digital content, including 3D assets, environments, characters, and animations, with minimal human input. These advancements are revolutionizing industries such as gaming, film, entertainment, architecture, and education by drastically reducing the time, cost, and labor required to produce rich, interactive experiences. The following sections explore these applications in greater detail, illustrating how businesses are leveraging these capabilities to push the boundaries of creativity and efficiency.

Game Development: Creating Dynamic Worlds and Realistic Characters

In the gaming industry, the application of Generative AI combined with Spatial Computing has fundamentally changed the development process, especially for VR and AR games. Traditionally, creating game assets such as landscapes, characters, and environments requires extensive manual work from artists and developers, taking months or even years. Generative AI now automates much of this process, using deep learning algorithms and GANs to design complex terrains, populate environments with flora and fauna, and create detailed characters in a fraction of the time.

For example, a game development studio might use AI to generate expansive landscapes that adjust based on gameplay mechanics, allowing players to experience varied terrains, such as deserts, forests, or urban cities, all dynamically rendered. AI can also create procedural content, meaning that every time a player enters a new area, the environment is freshly generated, making each gaming experience unique. In **RPG (role-playing game)** development, AI-driven character generation systems can produce a diverse array of **NPCs (non-player characters)** with distinct facial features, clothing, and behaviors, enriching the interactive experience without requiring manual design for each character.

Furthermore, Generative AI can animate these assets, creating lifelike movements and behaviors for characters. This is particularly valuable in VR, where immersion depends heavily on realistic interactions. AI-generated animations ensure that characters move fluidly and respond to user actions in a believable manner. This capability not only saves time but also enhances the depth and engagement of VR games by making the digital world feel more alive and responsive to players.

Film and Entertainment: Virtual Sets and Digital Doubles

The entertainment industry, particularly the gaming and film sectors, stands to benefit immensely from the integration of Generative AI and Spatial Computing. In gaming, developers can use these technologies to build complex game environments and characters that react dynamically to player actions. AI can generate new levels, quests, or scenarios based on how players interact with the game, providing an ever-evolving gaming experience that adapts to the preferences and skill levels of individual players. Such adaptive environments create a more immersive and engaging experience, keeping players invested for longer periods.

In film and television, the integration of AI and spatial technologies allows for interactive storytelling. Rather than consuming content passively, audiences can become active participants in a narrative. VR films and experiences generated using AI can provide branching storylines that change based on viewer choices, creating a new genre of immersive storytelling where each user experiences a unique version of the narrative. Production studios can further capitalize on this by offering episodic updates or interactive features that extend the lifecycle of content, fostering continuous viewer engagement.

Music and performing arts are also integrating these technologies. For instance, live concerts in VR, powered by AI-generated visual effects, can offer an unparalleled experience where the audience feels as if they are in the center of the action. The use of spatial audio and AI-generated light shows can synchronize with the rhythm and mood of the performance, creating a deeply immersive and personalized concert experience that can be enjoyed from anywhere in the world.

The film and entertainment industry is leveraging Generative AI in conjunction with Spatial Computing to transform traditional production processes. Virtual sets and digital doubles are two of the most significant innovations enabled by these technologies. Filmmakers can now create entire virtual environments—cities, landscapes, or fantastical worlds—without needing physical locations or elaborate set constructions. These AI-generated sets are highly detailed and dynamic, allowing directors and production teams to adjust scenes quickly and at a lower cost compared to building physical sets.

Digital doubles, or virtual representations of actors, are also becoming increasingly sophisticated through AI. By scanning an actor's appearance and movements, AI systems can generate realistic digital replicas that can perform stunts, interact with virtual environments, or even age or de-age characters for narrative purposes. This technology has proven particularly valuable for action and sci-fi films where complex or dangerous scenes can be simulated without putting actors at risk.

Moreover, digital doubles reduce the need for extensive post-production work, as AI can simulate lighting, camera angles, and movement in a virtual setting, allowing for real-time adjustments during filming.

The automation of background elements, such as crowd scenes, is another area where Generative AI shines. AI models can create thousands of distinct characters, each with unique appearances and behaviors, populating large digital spaces like stadiums, marketplaces, or battlefields without requiring extras or manual animation for each character. These AI-generated crowd simulations move and react realistically within scenes, significantly enhancing the visual scale and dynamism of movies and TV shows while minimizing production costs and logistics.

Generative AI's ability to autonomously produce high-quality digital assets like images, 3D models, videos, and animations aligns perfectly with the Apple Vision Pro's capabilities, offering ground-breaking tools for artists, designers, and developers (more on the Apple Vision Pro in the next chapter). Traditionally, content creation in AR/VR environments has required specialized skills and extensive resources, but the combination of Generative AI and the Vision Pro could democratize this process, enabling anyone to create immersive content with minimal technical expertise.

For instance, a designer wearing the Vision Pro could sketch a concept for a virtual space or an object, and the Generative AI could immediately transform that sketch into a fully rendered 3D model. The AI could then offer variations based on the designer's style or even optimize the design for various use cases, such as gaming, architecture, or marketing. This level of automation and assistance would accelerate the creative process, allowing professionals and hobbyists alike to produce high-quality digital assets in real time.

Architecture and Real Estate: Immersive Visualizations and Virtual Tours

Generative AI and Spatial Computing are revolutionizing the fields of architecture and real estate by automating the creation of 3D models and immersive environments. Architects can use AI to quickly generate building designs and interior layouts based on specific criteria such as sustainability, material efficiency, or client preferences. These AI-generated designs can be visualized within immersive AR or VR environments, allowing architects and clients to walk through virtual buildings before construction begins. This immersive visualization provides a more intuitive understanding of spatial layouts and aesthetics, improving decision-making and collaboration between stakeholders.

In the real estate sector, AI-generated virtual tours are becoming a powerful marketing tool. Real estate agencies can use Generative AI to create highly detailed virtual replicas of properties, enabling potential buyers to explore spaces remotely using VR headsets or AR applications. These AI-generated tours can simulate different lighting conditions, interior designs, or renovations, providing a personalized and engaging experience for clients. By automating the creation of these immersive environments, agencies can showcase multiple properties simultaneously, increasing efficiency and expanding their market reach beyond physical limitations.

Education and Training: Interactive Learning Environments

Generative AI is also playing a crucial role in transforming educational and training environments. By automating the creation of immersive simulations, AI enables the development of highly interactive and customizable learning experiences. For instance, in medical training, AI-generated VR simulations can create realistic operating rooms where medical students practice surgical procedures in a risk-free environment. These simulations are tailored based on real-world data, such as patient anatomy and surgical techniques, offering a level of hands-on training that is both immersive and educationally effective.

Similarly, in corporate training, businesses use AI-driven immersive environments to simulate real-world scenarios, such as customer interactions, machinery operation, or emergency response. Employees can engage in these simulations using VR headsets, allowing them to practice skills and decision-making in a controlled yet realistic setting. This not only enhances learning outcomes but also reduces training costs, as companies can provide scalable virtual environments without the need for physical resources or instructors. AI models continuously adapt these simulations based on user performance, ensuring that training experiences remain challenging and effective as employees improve their skills.

Educational institutions are also leveraging AI-generated content for interactive learning experiences. For example, history students can explore AI-generated reconstructions of ancient civilizations in VR, allowing them to walk through historical sites, witness events, and interact with AI-driven characters representing historical figures. This immersive approach not only engages students but also deepens their understanding by enabling experiential learning that goes beyond textbooks and traditional lectures. AI also allows for tracking students' progress, which it then follows up with specific study tracks that take into account any deficiencies found.

Marketing and Retail: Dynamic Product Showcases and AR Experiences

In marketing and retail, Generative AI combined with Spatial Computing is enhancing customer engagement by automating the creation of dynamic product showcases and interactive AR experiences. Brands can now generate digital replicas of their products and integrate them into AR apps that allow customers to visualize items in their own environment. For instance, furniture companies use AI-generated 3D models to let customers visualize how a sofa or table would look in their living room using a smartphone app. This application of AR enhances the shopping experience by providing customers with a realistic and personalized preview of products, increasing the likelihood of purchase and reducing return rates.

AI-driven AR experiences also enable interactive marketing campaigns that engage consumers in new ways. For example, brands can create virtual try-on experiences for clothing, accessories, or cosmetics, where AI-generated models of products adapt to the user's body or facial features in real time. These dynamic showcases not only offer a unique and engaging experience but also collect valuable user data, allowing businesses to tailor future marketing strategies and product recommendations based on customer preferences and interactions.

In addition, businesses are integrating AI-generated AR experiences into public spaces. Interactive billboards and digital installations powered by AI can display dynamic content that changes based on the audience's demographics or behavior. For example, an AR installation in a shopping mall might adjust its visuals and offers in real time, depending on the age group, gender, or interests of passersby. These targeted experiences create a more personalized and memorable interaction with the brand, enhancing marketing effectiveness and increasing customer loyalty.

Automotive Industry: Automated Prototyping and Virtual Testing

The automotive industry is utilizing Generative AI and Spatial Computing for automated prototyping and virtual testing, significantly reducing the time and cost of developing new vehicle models. AI systems can generate numerous design iterations based on parameters such as aerodynamics, material efficiency, and aesthetic appeal. These AI-generated models are then visualized and tested within VR environments, allowing engineers to assess the performance and safety of vehicles under various simulated conditions, such as different terrains, weather scenarios, or crash tests.

By using virtual testing environments, automotive companies can iterate designs more rapidly, identifying and addressing issues before creating physical prototypes. This not only streamlines the development process but also reduces material waste and resource expenditure. Moreover, these AI-driven simulations enable the testing of cutting-edge technologies, such as autonomous driving features, in a controlled virtual space before real-world implementation, improving safety and reliability.

AI-generated content also extends to marketing within the automotive sector. Brands are creating immersive AR experiences that allow potential customers to explore vehicle interiors, customize features, and visualize how cars would appear in different colors or with various accessories. These interactive showcases provide an engaging and informative experience, enhancing customer decision-making and supporting sales. On the production side, it's possible to create digital doubles of cars in order to stress test them, which could help to estimate car parts' durability and possible manufacture problems.

The ability of Generative AI to automatically generate content for immersive environments is revolutionizing industries by enhancing creativity, efficiency, and scalability. From game development and film production to architecture, education, and retail, businesses are leveraging AI's capabilities to create dynamic, interactive experiences that engage users and streamline operations. As these technologies continue to evolve, the potential for even more sophisticated and personalized content generation will expand, driving further innovation and transformation across sectors. However, as the use of AI-generated immersive environments grows, businesses must also navigate the challenges of ensuring ethical use, managing costs, and integrating complex systems to fully realize the benefits of this groundbreaking technology.

Now that we have talked about how AI could be put to use for a number of industries, we will turn to how AI and data work together.

Personalized Customer Experiences through Data-Driven Insights

The current digital era is characterized by heightened competition and shifting customer expectations, pushing businesses to innovate continuously to capture and retain customer engagement. The use of AI, combined with Spatial Computing, offers a groundbreaking approach to personalized customer experiences by harnessing the power of data-driven insights. This synergy allows brands to craft hyper-individualized customer journeys that dynamically adapt in real time, enhancing user engagement, satisfaction, and loyalty. These advanced technologies transform traditional methods, providing precision and relevance at a level previously unattainable.

AI's Comprehensive Role in Understanding Customer Behavior

AI's transformative potential in personalizing experiences begins with its ability to process and make sense of massive datasets derived from diverse customer touchpoints. By leveraging sophisticated algorithms, including ML and neural networks, AI models can extract meaningful insights from complex datasets that include purchase histories, browsing habits, social media interactions, and even contextual data such as geolocation and demographic information.

The power of AI lies in its capability to recognize intricate patterns and correlations that humans might overlook. For example, AI can not only identify when customers are more likely to make purchases but also predict which product categories they are inclined to choose based on historical data and contextual factors such as seasonality, promotional events, and even weather patterns. Beyond these basic patterns, AI models can analyze sentiment from customer reviews, assess their engagement with content on social media, and track unpurchased items frequently viewed in e-commerce stores.

By continually analyzing these data points, AI builds a dynamic and evolving profile of each customer. This profile is refined with every interaction, ensuring that the customer's preferences and needs are updated in real time. With such precision, businesses can anticipate and respond proactively to shifts in customer behavior, creating experiences that are not just relevant but genuinely tailored to the individual's evolving preferences.

Spatial Computing: Bridging Insights and Real-Time Experiences

AI's ability to derive insights is only half of the equation; the real challenge is implementing these insights in a way that feels seamless and engaging. Spatial Computing, including AR, VR, and MR, enables businesses to take the data-driven understanding generated by AI and translate it into immersive, interactive customer experiences. These technologies empower businesses to create digital environments that adapt instantly based on the AI's real-time analysis of customer behavior and preferences.

Consider the retail sector, where Spatial Computing is redefining the shopping experience. AR applications, for instance, can provide customers with tailored recommendations and interactive product showcases as they navigate a store. Imagine a scenario where, as a customer walks through an AR-enabled clothing store, the system accesses their previous shopping history and current preferences in real time. It can display virtual outfits that match the customer's style directly through their AR interface, suggest complementary accessories, and even provide exclusive promotions based on their past loyalty points. Such personalized touchpoints not only enhance the shopping journey but also significantly increase the likelihood of purchases by presenting options most relevant to the customer's taste and preferences.

Beyond the physical store, AR applications can extend this personalized approach to online shopping as well. Customers can use AR features on their smartphones to visualize products in their home environments before making a purchase, such as placing virtual furniture in their living room. The system, supported by AI, can adjust these suggestions based on past online behavior and external factors like trending design styles or seasonal sales, providing a customized and visually compelling shopping experience.

Revolutionizing Entertainment with AI-Driven Spatial Computing

The entertainment industry is another sector where AI and Spatial Computing are creating groundbreaking experiences. Streaming platforms, gaming companies, and immersive media providers are increasingly turning to Generative AI models to personalize entertainment on an unprecedented scale. These AI systems can design unique, interactive experiences tailored to each user's preferences, transforming how entertainment is consumed.

In the context of VR, imagine a platform that hosts live virtual concerts. The AI system, using its knowledge of the user's past behavior and musical preferences, can personalize every aspect of the event—from the setlist to the visual effects and stage environment. It could adjust the lighting or the visual scenery based on the genres and moods that resonate with the user. By creating a virtual concert experience that feels handcrafted for each individual, entertainment providers not only enhance user satisfaction but also increase the likelihood of customer retention and platform loyalty. Furthermore, such personalized experiences open up new revenue streams through subscription services, interactive content, and the sale of virtual merchandise exclusive to the user.

Creating Competitive Advantage through Immersive, Personalized Experiences

The integration of AI and Spatial Computing provides businesses with a powerful competitive advantage, enabling them to deliver experiences that traditional marketing or engagement methods cannot match. As these technologies evolve, companies can use the data collected from each interaction to continuously refine and enhance the user experience, making it increasingly tailored and relevant over time. The result is a dynamic and evolving ecosystem where businesses not only engage customers more effectively but also build long-lasting relationships based on value and relevance.

For example, within the hospitality industry, the impact of this technological integration is immense. Hotels can use AI-powered AR systems to enhance every aspect of the guest experience. Upon arrival, the hotel's AR system could offer a customized digital concierge that welcomes guests with an interactive guide featuring activities and experiences tailored to their interests and dietary preferences.

It might suggest nearby restaurants aligned with their tastes or show exclusive events happening during their stay. By making each guest's stay feel bespoke and accommodating, hospitality brands can foster a sense of personalized care, enhancing customer satisfaction and loyalty, ultimately increasing the likelihood of repeat visits.

As AI and Spatial Computing technologies evolve, their applications will proliferate across multiple industries, offering unprecedented opportunities for businesses to innovate and redefine customer experiences. Companies that invest in these technologies will not only enhance engagement but also build systems capable of adapting dynamically to customer needs.

In healthcare, for instance, AI and Spatial Computing can be used to create customized patient care plans. Imagine an AR application that guides patients through their treatment, providing real-time updates and visualizations based on their medical history, genetic information, and health status. This type of technology can transform patient care by offering visual and interactive support, ensuring that patients understand their treatment process and adhere to personalized health plans. Going further, AI could be used to predict potential diseases by analyzing a patient's history.

In education, the impact of AI and Spatial Computing could be equally transformative. Virtual classrooms could use AI to adapt learning materials and methods based on each student's learning style and performance. Students could interact with immersive content that adjusts in real time to their progress, helping educators identify learning gaps and providing immediate, personalized feedback. Such an approach not only enhances student engagement but also improves learning outcomes by tailoring the educational experience to meet individual needs.

The integration of AI with Spatial Computing offers businesses a unique opportunity to create personalized, immersive, and engaging customer experiences that adapt dynamically to each individual. By leveraging these technologies, companies can move beyond traditional engagement methods, providing real-time, responsive interactions that build long-term relationships and foster brand loyalty. As these technologies advance, their role in customer engagement strategies will only become more prominent, making it imperative for forward-thinking businesses to invest in these tools now. The ability to offer deeply customized, immersive experiences will not just be a competitive advantage but a necessity in the future digital economy.

We now talk about how digital twins can aid smart cities and urban development.

Smart Cities and Urban Development: Real-Time Digital Twins

The concept of smart cities is becoming increasingly prominent as urban areas seek sustainable, efficient, and technologically integrated solutions for growth and management. At the forefront of this revolution are Generative AI and Spatial Computing, which enable the creation and management of smart cities through the development of real-time digital twins. These digital twins act as comprehensive virtual replicas of physical environments—ranging from individual buildings and city blocks to entire metropolitan areas. They are not just static models; they integrate and update data continuously from IoT devices, sensors, and other data sources spread across the urban environment, offering an accurate and real-time representation of various city dynamics.

Understanding Digital Twins: A Complex Synthesis of Data and Spatial Computing

A digital twin is a sophisticated, dynamic model that mirrors the physical characteristics, activities, and environmental conditions of a real-world space. By integrating vast amounts of data from IoT devices, sensors, cameras, and other sources, digital twins provide a holistic view of city operations and conditions. These devices capture real-time information on diverse aspects such as traffic flow, air quality levels, energy and utility usage, and infrastructure integrity. For example, cameras can monitor traffic and pedestrian movement, while environmental sensors track pollution levels and temperature changes across different parts of the city.

The true power of digital twins lies in their ability to process and visualize this information through Spatial Computing. These virtual environments use AR, VR, and MR technologies to create immersive and interactive models that urban planners, businesses, and city managers can explore and manipulate. For instance, a planner can virtually navigate the city's streets, interact with infrastructure components like traffic lights and utility lines, and visualize the immediate impact of any changes they might implement.

Generative AI: Transforming Digital Twins into Intelligent Urban Systems

The integration of Generative AI into digital twins takes their capabilities beyond simple real-time visualization and monitoring. AI systems can analyze the immense volumes of data collected, identifying patterns, predicting outcomes, and simulating scenarios. By harnessing the predictive and analytical capabilities of AI, digital twins become intelligent systems capable of suggesting and implementing optimizations for city management in real time.

Optimizing Traffic Flow and Reducing Congestion

One of the most significant applications of AI-enhanced digital twins is in optimizing traffic management. Traditional traffic systems often rely on static models or basic, reactive approaches, which can lead to inefficiencies and congestion. By contrast, a digital twin equipped with AI can continuously analyze traffic data from sensors and cameras across the city. AI algorithms detect patterns in vehicle movement, identify congestion hotspots, and predict traffic surges based on time of day, weather conditions, and events like concerts or sports games.

City planners can use this information to implement dynamic traffic control measures. For example, AI can automatically adjust traffic light timings in real time to ease congestion or reroute vehicles through alternative paths to distribute traffic more evenly. Furthermore, the system can simulate the impact of new policies, such as the introduction of dedicated bike lanes or new public transport routes, allowing planners to visualize and assess their effectiveness before implementation.

Enhancing Utility Management and Infrastructure Maintenance

Digital twins also revolutionize the way cities manage utilities such as water, gas, electricity, and waste disposal. Utility companies can visualize their entire network in a virtual space, with real-time data showing usage patterns, operational statuses, and potential issues like leaks, outages, or system overloads. AI enhances this capability by analyzing usage trends, predicting maintenance needs, and optimizing distribution. AI could even be used together with weather models to predict droughts or floods, which would be helpful in optimizing water reserves.

For instance, an AI-powered digital twin can monitor water usage throughout a city, identifying areas with unusually high consumption that may indicate leaks or inefficient systems. By pinpointing these issues in real time, utility managers can prioritize maintenance efforts, dispatch repair teams to specific locations, and minimize service disruptions. Similarly, for electricity grids, AI can balance load distribution, optimize energy flow, and forecast demand based on patterns such as weather forecasts or historical usage data, ensuring efficient energy management and reducing the likelihood of blackouts.

Business Opportunities in the Digital Twin Ecosystem

The convergence of Generative AI, Spatial Computing, and digital twins presents vast opportunities for businesses, particularly those specializing in infrastructure, urban planning, technology integration, and smart devices. The market for digital twins is expanding as cities and municipalities increasingly adopt smart city initiatives to improve efficiency and sustainability. Companies that can offer digital twin solutions or complementary services stand to gain significantly.

Predictive Maintenance and Infrastructure Optimization

One major business application is predictive maintenance. Companies can leverage digital twins to offer predictive maintenance services for critical infrastructure components like bridges, roads, buildings, and utility lines. By analyzing data collected through sensors embedded in these structures, AI can forecast potential issues before they become significant problems, enabling proactive maintenance strategies. For example, an AI system might detect subtle shifts or vibrations in a bridge that indicate early signs of structural stress, prompting timely interventions that prevent catastrophic failure and costly repairs.

This type of service not only extends the lifespan of infrastructure but also enhances public safety, making it a valuable offering for both public and private sector entities. Infrastructure optimization can also include smart energy solutions where companies use digital twins to monitor and optimize energy usage across large complexes or neighborhoods. Businesses can propose energy-saving measures, such as installing smart grids or adaptive lighting systems, which adjust to the city's needs in real time based on occupancy data or daylight availability.

Urban Planning and Development Consultancies

The power of digital twins extends to urban planning and development consultancies, which can use these tools to offer smarter, more informed planning solutions to governments and developers. For instance, consultants can create virtual models of proposed developments, simulating how these new structures would interact with existing city elements. They can test different building materials, layouts, and energy solutions to determine the most efficient, sustainable, and visually appealing options.

This application is particularly beneficial in large-scale projects like smart city developments or urban revitalization initiatives. By using digital twins, planners can identify and mitigate potential issues—such as increased traffic congestion or reduced green space—before construction even begins. It allows for a collaborative approach where stakeholders, including architects, government officials, and community representatives, can interact with the virtual model, suggest changes, and see the immediate impact of those suggestions.

Transforming Public Safety and Emergency Response

Another critical application of real-time digital twins is in enhancing public safety and emergency response. By integrating AI and Spatial Computing, digital twins can provide city officials and first responders with an accurate, up-to-the-minute overview of an emergency situation. For example, in the case of a natural disaster like a flood or earthquake, a digital twin can simulate the spread of damage and predict its impact on infrastructure, guiding emergency services on where to deploy resources most effectively.

AI can also analyze past incidents to improve future response strategies. By studying patterns in emergencies such as traffic accidents or crime hotspots, AI systems can recommend preventive measures, such as optimizing patrol routes for law enforcement or identifying locations for new traffic safety features. Real-time monitoring during events like large public gatherings allows authorities to adjust their strategies based on crowd movement data, enhancing safety and minimizing risks.

The Ethical Considerations and Privacy Challenges of Smart Cities

While the benefits of digital twins and AI for urban development are substantial, they also raise important ethical and privacy considerations. The vast network of sensors and IoT devices used to collect data for digital twins can generate concerns about surveillance and data privacy. Cities must navigate these challenges carefully to ensure that data collection is transparent and secure.

Implementing stringent data privacy protocols and providing citizens with control over how their data is used are essential steps. Cities need to ensure that data collection is anonymized where possible and that only relevant information is gathered to maintain public trust. Regulations and policies must be developed to clearly communicate how data contributes to city management and to offer residents opt-in or opt-out options for specific services. By addressing these concerns proactively, smart cities can harness the power of digital twins while maintaining the trust and support of their citizens.

Looking Ahead: The Future of Smart Cities and Digital Twin Technology

The future of smart cities is undoubtedly tied to the development and evolution of digital twin technology. As AI systems grow more sophisticated and Spatial Computing technologies become more accessible, the capabilities of digital twins will expand, offering even more accurate, efficient, and interactive models of urban environments. Cities that invest in these technologies today will position themselves at the forefront of innovation, able to provide their residents with improved services, reduced environmental impact, and enhanced quality of life.

Future developments include the full integration of digital twins with 5G networks, enhancing real-time responsiveness and enabling even larger-scale simulations and data-processing capabilities. AI could also be used to predict long-term urban trends, such as population growth patterns or environmental changes, allowing for proactive planning and sustainable development initiatives.

In conclusion, the integration of Generative AI and Spatial Computing in the creation of digital twins is revolutionizing urban development. These technologies enable cities to become smarter, more efficient, and more responsive to the needs of their residents. As cities continue to grow and face new challenges, the adoption and refinement of digital twin technology will be essential in shaping the urban landscapes of the future, driving sustainable and intelligent development practices that benefit society as a whole.

Economic Impact

The economic impact of Generative AI and Spatial Computing is profound, fundamentally altering traditional business models, driving innovation, and creating entirely new revenue streams. This transformation spans industries, reshaping how businesses operate, engage with consumers, and generate value. The economic effects of these technologies are not just limited to cost savings and efficiency gains; they also catalyze new markets, redefine labor dynamics, and introduce unprecedented ways to monetize digital and virtual assets.

Automation and Cost Efficiency

One of the most immediate economic impacts of Generative AI is its ability to automate both repetitive and creative tasks, significantly reducing the reliance on human labor. In manufacturing, for example, AI can be integrated into design and production processes, automating everything from the generation of product designs to the execution of manufacturing tasks using AI-driven robotics. This level of automation minimizes errors, reduces production times, and lowers labor costs, leading to substantial savings for businesses.

In content-driven industries such as marketing, media, and entertainment, Generative AI automates the creation of visuals, text, and multimedia content, enabling businesses to produce high volumes of personalized and targeted material at a fraction of the cost of traditional methods. For instance, AI-generated advertising copy or social media content can be tailored to individual users based on demographic data and behavior patterns, enhancing marketing effectiveness while minimizing the costs associated with human copywriters and graphic designers.

Spatial Computing further amplifies these efficiencies through the use of digital simulations. For instance, in sectors like automotive or aerospace, creating and testing physical prototypes is an expensive and time-consuming process. With Spatial Computing, companies can build and test these prototypes in a virtual space using digital twins—accurate 3D representations of physical objects. This approach not only reduces the time and costs associated with physical prototyping but also allows for iterative design improvements to be implemented quickly and efficiently.

Enhanced Productivity and Streamlined Operations

Generative AI and Spatial Computing enable faster product cycles and streamlined operations across industries. In fields such as architecture and engineering, AI-generated designs are increasingly used to create complex structures and components that optimize efficiency and aesthetics. When these AI-driven designs are combined with Spatial Computing, professionals can visualize, test, and adjust models in immersive environments before moving to physical implementation. This integration accelerates the design process and reduces the likelihood of costly errors, leading to shorter project timelines and increased productivity.

In corporate environments, the use of Spatial Computing for virtual meetings and collaboration—often combined with AI tools—helps teams work more effectively, especially in a remote or hybrid setting. Digital twins of office spaces, factories, or even entire organizational workflows allow companies to monitor and optimize performance in real time, making adjustments based on AI-driven insights. This capability not only enhances operational efficiency but also supports decision-making with data-backed simulations, further contributing to profitability.

Emergence of New Business Models

The convergence of Generative AI and Spatial Computing has led to the creation of entirely new business models, expanding the economic landscape in ways previously unimagined. One of the most significant developments is the rise of **AR Experiences as a Service (ARaaS)**. Companies in various sectors, including retail, healthcare, and real estate, are subscribing to platforms that provide customizable and immersive AR experiences without the need for heavy investments in hardware or software infrastructure. For example, real estate firms are using ARaaS platforms to offer virtual property tours, allowing potential buyers to explore properties remotely with high levels of immersion and interactivity.

Subscription-based models are also emerging in entertainment and education, where VR and AR platforms provide immersive learning environments or personalized entertainment experiences on a monthly or annual basis. These services leverage AI to tailor content based on user preferences and engagement patterns, enhancing user retention and opening new streams of recurring revenue for businesses.

In the marketing sector, businesses are investing in AI-driven AR campaigns that engage consumers directly through their mobile devices. For example, interactive AR experiences in public spaces or through social media can allow users to interact with virtual products, try on digital clothing, or play branded games.

Fashion brands, for instance, are selling virtual clothing that can be worn by avatars in AR and VR environments. These digital outfits, often called “skins,” have become popular in gaming platforms where players pay for customization options for their characters. For instance, a company may offer a subscription service where users can attend virtual concerts, visit virtual museums, or engage in interactive narratives that are constantly updated by AI to keep the experience fresh. Companies pay for these immersive campaigns on a per-interaction or subscription basis, transforming traditional marketing budgets into dynamic investments with measurable returns.

Digital Assets and Virtual Economies

The ability to create and manage digital assets like virtual clothing, real estate, and entire virtual worlds has opened new avenues for monetization. In the fashion industry, brands are designing digital apparel that can be worn by avatars in VR and AR environments. This trend, often referred to as “digital fashion,” not only allows brands to reach new audiences in virtual spaces but also taps into emerging digital economies where consumers are willing to pay for exclusive, AI-generated virtual items. Companies are capitalizing on this by offering limited-edition digital products, creating scarcity and demand similar to that seen with physical products.

The gaming industry has perhaps seen the most dramatic impact from this shift. In-game economies, where players purchase AI-generated virtual items such as skins, weapons, and accessories, have grown into a multi-billion dollar marketplace. AI-generated assets can be customized in real time, offering endless variations that enhance the player’s experience and sense of individuality. The integration of Spatial Computing adds another layer, enabling players to interact with these digital items in immersive 3D environments, further enhancing their value and appeal.

The concept of virtual real estate is another lucrative frontier. Companies and developers are investing in digital plots within virtual worlds or metaverses, creating spaces that are used for a variety of commercial purposes, including retail storefronts, event venues, and social hubs. Businesses are buying, developing, and leasing these virtual properties, creating new economic models based entirely on digital real estate management. Digital real estate platforms allow for transactions using cryptocurrency, turning virtual property trading into a profitable business with its own financial ecosystem.

Economic Impact on Traditional Industries

The economic influence of Generative AI and Spatial Computing is not confined to digital-native companies; it also extends to traditional industries. In construction, for example, Spatial Computing allows for the visualization of construction sites and progress through AR overlays, providing real-time data that improves decision-making and reduces construction delays.

Generative AI assists by optimizing building designs for sustainability and cost-efficiency, ensuring that projects are both environmentally friendly and economically viable.

In healthcare, AI-powered virtual assistants and Spatial Computing tools are used for telemedicine, remote diagnostics, and surgical training. These innovations reduce the need for physical resources and labor, allowing healthcare facilities to deliver services more efficiently while minimizing costs. The ability to simulate medical procedures in VR environments not only enhances training but also reduces the risk and expense associated with hands-on learning.

Education is another traditional sector being transformed economically. With the rise of AI-powered learning platforms and VR classrooms, institutions can provide interactive and personalized learning experiences at scale. Students can attend virtual classes or explore immersive historical, scientific, and cultural experiences that would be impossible to replicate physically. These solutions offer cost-effective and scalable models for educational institutions, creating new revenue opportunities through digital course offerings and global student enrollment.

Long-Term Economic Trends and Future Projections

As the integration of Generative AI and Spatial Computing deepens, several long-term economic trends are emerging. The ongoing shift toward digital economies is likely to accelerate, with businesses increasingly focusing on monetizing virtual and digital assets. The rise of the metaverse, where immersive digital spaces serve as platforms for commerce, social interaction, and entertainment, is expected to create entire new economic ecosystems that parallel traditional physical economies.

The decline in the costs associated with these technologies, driven by advancements in AI algorithms and hardware miniaturization, will make them increasingly accessible to **small and medium-sized enterprises (SMEs)**. As SMEs gain access to these capabilities, the economic impact will become even more widespread, leading to a democratization of the benefits associated with Generative AI and Spatial Computing.

Additionally, the job market is likely to evolve significantly. While automation will undoubtedly replace certain roles, especially those focused on repetitive tasks, new job categories will emerge. The demand for specialists in AI ethics, digital asset management, VR development, and AI-enhanced design will increase, creating new opportunities for education and workforce development. Businesses that invest in upskilling employees to adapt to these new technologies will not only mitigate job displacement but also benefit from a more tech-savvy workforce capable of driving further innovation.

Challenges and Considerations

The integration of Generative AI and Spatial Computing offers immense business potential, but it also brings a complex set of challenges that need careful consideration and strategic management. These challenges span ethical, technical, financial, and regulatory dimensions. Addressing them is crucial for businesses to fully capitalize on these technologies while minimizing risks and ensuring long-term success. This section delves into these challenges in detail, providing an expanded understanding of the key issues businesses face and the measures required to overcome them.

Ethical and Privacy Concerns

One of the most pressing challenges in the deployment of Generative AI and Spatial Computing is ensuring the ethical use of data and AI-generated content. These technologies are inherently data-driven, often relying on vast amounts of personal and behavioral information to create personalized experiences or generate content. This raises significant concerns about privacy and the security of sensitive data. Businesses must navigate a complex web of regulations, like the **General Data Protection Regulation (GDPR)** in Europe, the **California Consumer Privacy Act (CCPA)**, and other emerging privacy laws worldwide.

Privacy risks arise when Generative AI models, which require extensive training data, access personal user information. This data could include location history, biometric information, or consumption patterns, particularly in applications where AI is combined with AR/VR interfaces that track user movement and interactions. Companies must implement robust data protection measures such as encryption, anonymization, and user consent protocols to protect user privacy. Failure to do so not only will expose businesses to regulatory penalties but also risks damaging their reputation and eroding consumer trust.

In addition to privacy, the ethical implications of AI-generated content must be considered. As AI systems become capable of creating highly realistic visual, audio, and textual content, the risk of misuse, such as deepfakes or misleading advertisements, increases. Companies need to establish ethical guidelines and transparency in AI usage to avoid misuse and ensure that the generated content aligns with societal norms and values. Developing frameworks for accountability, transparency, and ethical AI use will be critical to maintaining public trust and avoiding legal complications.

High Costs and Accessibility Barriers

The adoption of Generative AI and Spatial Computing technologies, particularly at scale, involves significant upfront investment. These costs are associated with acquiring advanced hardware, such as AR/VR headsets, powerful computing infrastructure for AI processing, and high-quality sensors necessary for spatial mapping and interaction. Furthermore, software development for these immersive environments is resource-intensive, requiring specialized expertise in AI, computer graphics, and spatial engineering.

For SMEs, these high costs can be prohibitive, limiting their ability to access and deploy these technologies effectively. This creates a potential divide where only larger corporations with significant capital and technical capabilities can afford to implement and benefit from these advanced solutions. Such a disparity could widen the gap between industry leaders and smaller players, ultimately affecting competition and innovation.

To address these cost barriers, businesses and technology providers must explore scalable and accessible solutions. Cloud-based AI and AR/VR platforms, for example, offer opportunities to reduce upfront costs by providing software and computational power as a service. Additionally, public-private partnerships and government subsidies could play a role in democratizing access, enabling SMEs to experiment with and adopt these technologies at a lower financial risk.

Technical Complexity and Integration Challenges

The integration of Generative AI and Spatial Computing is a highly complex task that requires the seamless connection of multiple software systems, AI models, and hardware devices. Achieving this level of integration demands sophisticated technical expertise, particularly when it comes to ensuring interoperability and scalability. AI models must be compatible with Spatial Computing platforms, and data must flow smoothly across different systems to provide real-time responsiveness and accuracy in AR/VR environments.

For businesses aiming to deploy these technologies, the technical challenge extends beyond simply acquiring the right tools; it involves developing or hiring specialized talent capable of integrating these systems effectively. There is a growing demand for experts in fields like AI engineering, Spatial Computing development, and cross-platform software engineering. However, the current talent pool is limited, leading to a skills gap that further complicates the adoption process.

Moreover, ensuring scalability is another critical consideration. Businesses need to design systems that can grow with increasing user demand without compromising on performance. In practical terms, this means investing in robust cloud infrastructure that can handle the high computational loads of Generative AI and Spatial Computing applications. Additionally, companies must implement systems that can support future software updates and hardware improvements without requiring complete overhauls. Addressing these technical challenges requires a long-term commitment to research and development, as well as partnerships with technology providers that can offer the necessary support and expertise.

Intellectual Property and Copyright Issues

Generative AI technologies have revolutionized creativity and automation, allowing for the creation of original content across various domains, including visual art, music, text, 3D models, and virtual environments. While this development expands possibilities for innovation, it also introduces significant legal and ethical complexities concerning **intellectual property (IP)** and copyright. Traditional copyright laws were designed with human creators in mind, clearly outlining the rights of individuals or organizations producing creative works. With AI systems now capable of generating content autonomously or in collaboration with humans, determining ownership rights becomes increasingly ambiguous and legally challenging.

The Ambiguity of Ownership: Who Owns AI-Generated Content?

A fundamental challenge with AI-generated content is the issue of ownership. When humans create a work—such as a painting, a song, or a novel—copyright laws typically recognize them as the rightful owners of the work. However, when AI systems produce art, designs, or other creative outputs, assigning ownership becomes unclear. This uncertainty raises several perspectives regarding who might rightfully claim ownership.

The developer who designs and builds the AI system could argue that, since the AI's creative capabilities are a direct result of their programming and algorithmic design, they hold the rights to the content generated. This perspective emphasizes that without their expertise and intellectual input, the AI would not be able to create.

Conversely, the company that deploys and uses AI for commercial purposes may assert that it owns the rights to the AI-generated content. When a business invests in AI technology, directs its use, and derives economic benefits from the outputs, it may claim that the rights belong to them as they oversee and manage the creative process, even if indirectly.

The end-user's involvement also adds another layer to the debate. Users who interact with the AI system, inputting specific prompts or guiding the process through iterations to achieve the final design, may feel entitled to ownership. They might view themselves as co-creators, having influenced the AI's output with their own creative intent and direction.

These varying perspectives illustrate the complexity of the issue, creating a "gray area" not adequately covered by existing copyright frameworks. As AI technology evolves and becomes more widespread, it is essential for legal systems, developers, and businesses to work together to establish guidelines that clearly define ownership and protect the interests of all parties involved.

Legal and Operational Risks for Businesses Utilizing Generative AI

The ambiguous nature of IP ownership in AI-generated content presents significant legal and operational risks for businesses. Companies using AI for generating content—whether for marketing campaigns, product design, branding elements, or digital assets—must be cautious of potential IP disputes and copyright infringement claims. Without clearly defined ownership, businesses could inadvertently use or distribute materials that may later be contested by developers, end-users, or other stakeholders. This uncertainty can lead to legal battles, financial penalties, and reputational damage if companies are found to be infringing on copyright laws.

To mitigate these risks, businesses must proactively address these complexities. Establishing clear guidelines and legal frameworks is crucial to determining ownership and usage rights for AI-generated content. Collaborating with legal experts to develop contracts, policies, and usage agreements that clearly state the terms of ownership and define how the AI outputs are managed under copyright laws can help businesses safeguard their interests. These agreements need to cover various scenarios, including the rights of the AI developers, the entities utilizing the technology, and the role of end-users in co-creating or influencing the output.

Authenticity, Traceability, and the Role of Blockchain Technology

Beyond the question of ownership, companies also face challenges related to the authenticity and traceability of AI-generated content. As digital art and virtual products become more prevalent, especially in sectors like gaming, digital fashion, and the virtual goods market, ensuring the authenticity and uniqueness of AI-generated content becomes critical. In a world where AI can produce nearly limitless variations of an image or design, distinguishing original work from derivative or replicated content becomes increasingly complex.

The emergence of blockchain technology offers potential solutions to these challenges. Blockchain's decentralized and immutable nature makes it an ideal tool for verifying the authenticity and ownership of digital assets. By integrating blockchain-based verification systems, businesses can create a secure record for AI-generated content, ensuring transparency in its creation, distribution, and ownership. This technology provides a way to manage IP rights effectively in digital and virtual environments by creating a traceable digital signature for each piece of content. Blockchain can also support smart contracts that automatically enforce ownership terms, royalties, and usage rights, providing creators and owners with fair recognition and compensation.

Incorporating blockchain into the management of AI-generated content could transform how businesses navigate IP challenges, offering a reliable and transparent way to authenticate and protect digital creations. By leveraging such technology, companies can build greater trust with consumers and partners, ensuring that their AI-generated assets remain secure and that all parties involved receive appropriate credit and compensation.

Establishing a Legal Framework for AI-Generated Content

For the successful adoption of AI-generated content, it is essential for regulatory bodies and industry leaders to develop comprehensive legal frameworks. Such frameworks should address the complexities of ownership, usage rights, and copyright management in the context of AI systems. This process involves reviewing existing copyright laws and adapting them to accommodate scenarios where non-human entities play a significant role in the creative process.

Policymakers, AI developers, businesses, and legal experts must collaborate to establish standards that clarify the roles and rights of each party involved in AI-generated content creation. Clear guidelines should outline the extent of ownership for AI developers, companies utilizing the technology, and users interacting with the system. The framework should also account for cases where multiple entities contribute to the creation process, establishing shared or co-ownership models where appropriate.

Additionally, it is crucial to build policies that ensure AI-generated content adheres to ethical standards and respects existing copyrights, especially when AI is used to remix, reference, or emulate the work of other artists. AI systems trained on vast datasets often incorporate elements from copyrighted materials, raising questions about the originality and legality of the outputs they produce. A well-defined legal framework should provide mechanisms for transparency, enabling creators to understand how their work is being used in training AI models and giving them the option to opt in or receive compensation.

The Path Forward: Balancing Innovation with Legal Certainty

The future of AI-generated content lies in balancing technological innovation with robust legal protections that provide clarity and security for creators, businesses, and consumers. As the use of Generative AI continues to expand, establishing comprehensive legal frameworks will be essential to fostering a safe and sustainable creative economy where all participants—developers, businesses, and users—understand their rights and responsibilities.

The continued evolution of these technologies will require ongoing collaboration between industries, legal institutions, and governments to keep pace with advancements and address emerging challenges. By developing adaptive and forward-thinking regulations, stakeholders can ensure that the creative potential of AI is harnessed responsibly, unlocking new opportunities while minimizing legal risks and protecting IP rights. As AI becomes more integrated into creative processes, it is vital that policies evolve alongside it, providing the clarity needed for businesses and individuals to navigate this complex and rapidly changing landscape effectively.

Regulatory Compliance and Global Standards

As Generative AI and Spatial Computing technologies gain traction and become integral to various industries, governments and regulatory bodies around the world are increasingly focusing on establishing comprehensive guidelines and regulations to manage their deployment and usage. The rise of these regulations is driven by concerns over privacy, data security, ethics, and the potential societal impacts of these advanced technologies. However, the regulatory landscape is far from uniform, with considerable variation across regions. This diversity in regulations creates significant challenges for businesses operating in multiple countries, as they must navigate a complex web of compliance requirements that often differ substantially from one jurisdiction to another.

The Regional Disparities in Regulatory Approaches

One of the most significant challenges global businesses face is the variation in regulatory approaches across different regions. In Europe, for example, GDPR imposes stringent requirements on how companies collect, process, and store personal data. GDPR mandates that businesses obtain explicit consent from users before using their data, offer options for individuals to access and delete their data, and ensure robust protection measures are in place to safeguard privacy. For companies deploying AI and Spatial Computing technologies in Europe, compliance with GDPR means implementing strict data-handling practices and transparency measures.

On the other hand, regions like the United States and parts of Asia may have less comprehensive data protection regulations, resulting in a regulatory environment that is comparatively lenient or fragmented. In the United States, regulations around data privacy and AI vary from state to state, with some states like California enforcing stricter rules (e.g., CCPA), while others have minimal requirements. Similarly, in Asia, regulations range from highly regulated markets like Singapore to less structured frameworks in developing regions.

This inconsistency forces multinational businesses to develop region-specific compliance strategies. Companies must tailor their data privacy practices, AI deployments, and Spatial Computing applications to align with the local laws of each market they operate in. Such localization efforts can be resource-intensive, requiring significant investments in legal expertise, technology adaptations, and operational changes to ensure compliance. The complexity of managing multiple, sometimes conflicting, regulatory requirements makes it essential for businesses to establish a comprehensive and flexible compliance framework capable of adapting to each jurisdiction's specific rules.

The Rise of AI Ethics Regulations and the Challenge of Responsible AI

Beyond privacy regulations, new regulatory measures are emerging to address the ethical implications of AI and its responsible use. As AI becomes more pervasive in decision-making processes and content creation, regulatory bodies are introducing guidelines aimed at ensuring transparency, fairness, and accountability. For instance, the European Union's AI Act sets out to categorize AI systems based on risk levels—ranging from minimal risk (such as spam filters) to high-risk applications (such as biometric identification systems or AI in healthcare). High-risk systems are subject to stricter requirements, such as proving the system's transparency, demonstrating robust data governance practices, and conducting thorough assessments to prevent bias or discrimination.

These emerging regulations place additional compliance demands on businesses using AI technologies. Companies must not only ensure that their AI systems comply with privacy laws but also demonstrate that their AI models operate transparently and ethically. This may involve disclosing how data is collected and processed, providing explainable AI models that reveal decision-making processes, and proving that their AI outputs are free from bias, misinformation, or harmful content. For businesses developing AI-driven content through Generative AI or using Spatial Computing to simulate and analyze real-world scenarios, these requirements add a significant layer of complexity.

To comply with these ethical regulations, companies need to develop comprehensive AI governance programs. This includes implementing processes for auditing and monitoring AI systems, ensuring that data inputs and model outputs are transparent and unbiased, and creating mechanisms to manage and mitigate risks associated with AI-generated content. Businesses may also need to engage in AI ethics training for their teams and establish cross-functional committees dedicated to overseeing compliance efforts. Investing in research and development to create explainable AI models that provide insights into how algorithms reach their conclusions can help businesses meet regulatory demands and build trust with stakeholders.

The Role of Compliance Programs and Real-Time Monitoring

Given the multifaceted nature of regulations surrounding AI and Spatial Computing technologies, businesses must build robust compliance programs capable of monitoring and managing their adherence to local, national, and international standards. These programs should integrate compliance considerations into every stage of product development and deployment, ensuring that legal and ethical requirements are addressed from the outset. For instance, incorporating privacy-by-design principles into software development processes can help companies create AI systems that are inherently compliant with data protection laws.

Real-time monitoring and reporting are critical components of these compliance programs. With regulations constantly evolving, especially in the rapidly advancing field of AI, businesses need tools that can track regulatory changes, assess their impact, and adjust compliance strategies accordingly. This proactive approach ensures that businesses remain compliant with new laws and guidelines, avoiding the risk of penalties or legal disputes.

Investing in **regulatory technology (RegTech)** solutions offers a viable pathway for businesses to manage these challenges efficiently. RegTech platforms provide automated compliance monitoring and reporting, enabling companies to keep track of regulatory requirements in different regions. These systems can analyze data flows, detect non-compliance issues in real time, and automatically generate reports for regulatory bodies. By integrating RegTech into their operations, businesses can reduce the manual effort and complexity associated with managing compliance, while also improving the accuracy and responsiveness of their compliance efforts.

The Consequences of Non-Compliance: Financial and Reputational Risks

The penalties for failing to comply with AI and data privacy regulations can be severe, impacting both the financial stability and reputation of a business. Under GDPR, for example, companies can face fines of up to €20 million or 4% of their global annual revenue, whichever is higher. Such fines can be crippling, particularly for SMEs that may not have the resources to absorb such costs. In addition to financial penalties, companies that fail to comply with regulations may also experience damage to their reputation, resulting in lost consumer trust, diminished brand loyalty, and negative media coverage.

Reputational damage can have long-term consequences, affecting a company's market position and ability to attract new customers or partners. In industries where trust and transparency are critical—such as healthcare, finance, and technology—compliance failures can lead to significant customer attrition and difficulties in securing business partnerships. Therefore, proactively engaging with regulatory bodies and demonstrating a commitment to ethical and compliant practices is essential for businesses to maintain their reputation and ensure sustained growth.

Proactive Engagement with Regulatory Bodies and Evolving Legislation

To effectively manage regulatory challenges, businesses must adopt a proactive approach, engaging with regulatory bodies and staying informed about legislative developments. This engagement is crucial, as early involvement in regulatory discussions allows businesses to provide feedback and insights that may influence the development of new rules. By participating in industry coalitions and advocacy groups, businesses can collaborate with other organizations to establish best practices and contribute to shaping regulations that are realistic and manageable.

Remaining informed about evolving legislation is also vital. As governments around the world continue to refine their approach to regulating AI and Spatial Computing, businesses must stay up to date with new developments to ensure compliance. This may involve establishing internal compliance teams tasked with monitoring legislative changes and evaluating their impact on the company's operations. Alternatively, businesses can partner with legal experts specializing in AI and technology law to receive timely advice and guidance.

Furthermore, companies must integrate compliance considerations into their product development and business strategies. By aligning product development with regulatory requirements from the start, businesses can minimize costly redesigns and reduce the risk of non-compliance. Compliance efforts should extend beyond legal obligations to include a commitment to ethical and socially responsible use of AI technologies, as this builds consumer trust and strengthens brand value.

Investing in RegTech Solutions: Automating Compliance for Efficiency

The use of RegTech solutions offers an efficient and scalable approach to managing compliance across different regions and industries. RegTech platforms leverage AI and automation to provide real-time monitoring, ensuring that businesses stay compliant with evolving regulations without the need for extensive manual oversight. These platforms can automatically assess whether a company's AI and Spatial Computing systems adhere to privacy, ethical, and data governance requirements, flagging potential non-compliance issues for review.

By implementing RegTech, businesses can streamline their compliance processes, reducing the burden on their legal and compliance teams. Automated systems can generate compliance reports, track the flow of sensitive data, and integrate with AI governance frameworks to provide a comprehensive view of a company's compliance status. Additionally, RegTech solutions can be customized to accommodate region-specific regulations, making them adaptable tools for multinational corporations facing diverse compliance challenges.

The Future of Compliance: Toward Harmonization and Global Standards

Looking forward, the global regulatory environment for AI and Spatial Computing technologies is likely to evolve toward greater harmonization. As the technologies become more widespread and the potential risks become clearer, international organizations and governments may work together to establish unified standards that simplify compliance for businesses operating across borders. Initiatives led by entities like the European Union, the United Nations, and other global regulatory bodies could pave the way for common frameworks that promote consistent and ethical AI practices worldwide.

In conclusion, navigating regulatory compliance in the context of Generative AI and Spatial Computing requires a proactive and multifaceted approach. Businesses must adapt to a complex and evolving landscape by developing robust compliance programs, engaging with regulatory bodies, and leveraging technology solutions like RegTech to automate and streamline their efforts. By taking these steps, companies can not only mitigate risks but also position themselves as leaders in responsible AI deployment, building trust with consumers, partners, and regulators alike. As global standards continue to emerge, businesses that proactively engage in shaping and adhering to these standards will be best positioned for sustainable growth and success in the AI-driven economy.

Environmental Impact and Sustainability

The integration of Generative AI and Spatial Computing technologies is transforming industries, enabling innovations in entertainment, business, urban planning, and other sectors. However, this technological advancement comes with a significant environmental footprint. The energy consumption required to power AI systems, train complex models, and support immersive experiences in AR/VR platforms is substantial, and much of this energy is currently sourced from non-renewable resources. The environmental sustainability of these technologies is, therefore, a growing concern as businesses and societies increasingly depend on them.

The Environmental Costs of AI and Spatial Computing

Generative AI models, particularly those based on large neural networks, demand immense computational power. Training these models requires massive data centers equipped with high-performance computing resources, such as GPUs and specialized AI chips. These data centers, in turn, consume vast amounts of electricity, contributing significantly to carbon emissions. A single training cycle for a state-of-the-art AI model can produce as much carbon dioxide as multiple transatlantic flights. The electricity used in these data centers often comes from fossil fuels, which intensifies the carbon footprint of AI development (although the Chinese app DeepSeek comes with claims that it uses much less resources than others, particularly in the United States).

Similarly, the maintenance of immersive environments in AR/VR platforms also requires substantial energy. These platforms rely on complex rendering processes and real-time data processing to deliver seamless experiences, all of which necessitate high computational capabilities. The ongoing operation of these systems, especially as AR/VR technologies scale and become more mainstream, further increases energy demands and carbon emissions.

In addition to the direct energy consumption, the lifecycle of hardware used for AI and Spatial Computing also impacts the environment. High-performance computing components, such as GPUs and data center servers, require energy-intensive manufacturing processes and create electronic waste when they become obsolete. The rapid pace of technological advancement in AI and Spatial Computing means that hardware turnover rates are high, adding to the growing e-waste problem.

Mitigating the Environmental Impact: Strategies and Solutions

As businesses adopt and scale AI and Spatial Computing technologies, it is crucial to address their environmental impact through sustainable practices. Companies must consider both immediate and long-term strategies to reduce energy consumption, minimize emissions, and promote sustainability throughout their operations.

Adopting energy-efficient AI models is a critical first step. Researchers and developers are exploring techniques such as model pruning, quantization, and knowledge distillation, which optimize models to achieve similar performance with significantly reduced computational resources. These techniques help in creating “greener” AI systems that consume less power during both the training and inference phases. For instance, lightweight models that are smaller in size and use optimized algorithms can reduce the overall carbon footprint without compromising on performance.

Investing in renewable energy sources for powering data centers is another essential strategy for businesses looking to mitigate their environmental impact. By transitioning to green energy sources like solar, wind, or hydroelectric power, companies can significantly reduce the carbon emissions associated with their operations. Some tech giants, such as Google and Microsoft, have already made substantial investments in renewable energy to power their data centers, setting an example for other businesses in the industry. Smaller companies can also partner with green energy providers or invest in energy certificates that guarantee their data consumption supports sustainable sources.

Decentralized and edge computing technologies offer further opportunities for reducing energy demands on central servers. Rather than relying solely on large, centralized data centers, businesses can distribute computational loads across smaller, localized edge servers. These edge servers are often situated closer to the end-users, reducing the need for long-distance data transmission and lowering latency. By processing data locally, edge computing minimizes the energy consumption associated with transferring and storing large datasets in centralized locations, making the overall system more energy efficient. This approach also enhances the scalability and responsiveness of Spatial Computing applications, providing a win-win for sustainability and user experience.

Incorporating Sustainable Practices into Business Models

The growing awareness of climate change and environmental sustainability is increasingly influencing consumer choices and behavior. Businesses that prioritize green technology solutions and actively demonstrate their commitment to sustainability are likely to gain a competitive advantage in the market. By integrating sustainable practices into their core operations, companies can attract environmentally conscious customers and investors who value eco-friendly brands.

One way that businesses can integrate sustainability into their operations is by conducting regular environmental impact assessments. These assessments involve evaluating the carbon footprint of their AI and Spatial Computing deployments, identifying areas where emissions can be reduced, and setting measurable goals for improvement. For example, businesses can track their energy consumption across data centers, measure emissions from hardware manufacturing, and monitor the lifecycle of devices used in their operations. With this data, they can make informed decisions about where to invest in green technologies and how to optimize their systems for efficiency.

Incorporating sustainability into the product development cycle is also essential. Businesses can design products and services that prioritize energy efficiency, such as optimizing AR/VR platforms to operate effectively on less powerful hardware or developing AI applications that can run efficiently on mobile and edge devices. In addition, companies can explore partnerships with eco-friendly hardware manufacturers who use sustainable materials and processes, reducing the overall environmental impact of their technology stack.

Moreover, businesses can embrace circular economy principles by developing strategies for managing electronic waste. Recycling and refurbishing programs can be established to extend the lifespan of hardware components, and partnerships can be formed with e-waste recycling firms to properly dispose of outdated or damaged equipment. Offering customers trade-in programs or incentives for returning used devices can also help reduce the amount of electronic waste generated, promoting a more sustainable lifecycle for technology products.

Leveraging Green Technology as a Competitive Advantage

The growing emphasis on environmental sustainability is not only a moral obligation but also a strategic opportunity for businesses. Companies that adopt green technology practices and showcase their efforts to reduce their environmental impact can enhance their brand reputation and build trust with eco-conscious consumers and stakeholders. In sectors like technology, manufacturing, and retail, where environmental issues are gaining prominence, demonstrating a commitment to sustainability can differentiate a company from its competitors.

Businesses that take proactive steps to reduce their carbon footprint can communicate their green initiatives through transparent sustainability reports. These reports highlight energy-saving measures, emissions reductions, and the use of renewable energy sources, providing stakeholders with a clear picture of the company's sustainability efforts. Companies that can prove their commitment to reducing emissions and promoting eco-friendly practices may attract investors who prioritize **environmental, social, and governance (ESG)** criteria when making investment decisions.

Furthermore, green technology initiatives can be leveraged as marketing tools to appeal to consumers who are increasingly making purchase decisions based on environmental values. Businesses can engage with customers through campaigns that highlight their sustainable practices, such as using renewable energy or developing energy-efficient products (and AI could even be used to think about the alternatives of having a greener environment in the future). This engagement not only enhances customer loyalty but also positions the company as a responsible and forward-thinking brand that aligns with the values of its consumer base.

The Role of Policy and Industry Collaboration in Promoting Sustainability

To further enhance the sustainability of AI and Spatial Computing, businesses must also advocate for policies and collaborate with industry stakeholders to promote environmental responsibility on a larger scale. Governments and regulatory bodies are increasingly recognizing the importance of sustainability, and new regulations focused on carbon emissions, energy efficiency, and electronic waste management are emerging. Companies should actively engage with these policymakers to help shape regulations that are both effective and practical, ensuring that industry needs are considered while promoting sustainable development.

Industry collaboration is also essential in developing and implementing green technology standards. Businesses, technology developers, and environmental organizations can work together to establish guidelines for energy-efficient AI practices, responsible e-waste management, and the development of sustainable hardware solutions. By pooling resources and expertise, the industry can accelerate the adoption of eco-friendly technologies and share best practices that benefit the entire sector.

Initiatives such as the AI for Earth program or the **Renewable Energy Buyers Alliance (REBA)** demonstrate how industry collaborations can drive meaningful progress in sustainability. These programs bring together companies, NGOs, and technology developers to support the use of AI and renewable energy in environmental protection projects, showcasing the potential for AI to be a force for good when aligned with sustainability goals.

The Future of AI and Spatial Computing in an Eco-Conscious World

As AI and Spatial Computing technologies continue to evolve, their environmental impact will remain a critical concern for businesses and policymakers alike. The future of these technologies will depend on the ability of companies to balance innovation with environmental responsibility, ensuring that their products and services are sustainable and efficient. The shift toward green AI practices, renewable energy integration, and decentralized computing solutions represents a vital step in this direction, but it will require ongoing commitment and investment from businesses to realize these goals.

As the public and regulatory pressure for sustainability intensifies, companies that prioritize eco-friendly practices in their AI and Spatial Computing deployments will not only contribute to environmental protection but also secure a more competitive and resilient position in the market. The alignment of technological advancement with sustainability will be a defining factor for successful businesses in the years to come, demonstrating that innovation and environmental responsibility can and must go hand in hand. By leading the way in sustainable technology practices, companies can build a future where AI and Spatial Computing not only drive progress but do so in a manner that supports the health of our planet.

Workforce Adaptation and Skill Gaps

The adoption of Generative AI and Spatial Computing necessitates a skilled workforce capable of managing and developing these technologies. However, the rapid evolution of these fields has led to significant skill gaps in the labor market. Many businesses struggle to find professionals with expertise in AI, Spatial Computing, and interdisciplinary fields that combine elements of computer science, engineering, and design.

Workforce adaptation is essential for businesses to remain competitive in a technology-driven economy. Companies must invest in upskilling and reskilling programs to bridge these gaps, offering training that equips employees with the skills needed to work with AI and Spatial Computing tools. Collaboration with educational institutions, industry groups, and online learning platforms can help businesses create tailored training programs that align with their strategic needs.

Additionally, businesses must navigate the cultural shift that comes with technological transformation. Integrating these advanced technologies often requires changes in business processes, organizational structures, and communication practices. Change management strategies, which include transparent communication, leadership buy-in, and employee involvement, are crucial to facilitate this transition and ensure that employees are not only prepared but also motivated to adapt to new technologies.

The challenges and considerations associated with Generative AI and Spatial Computing are multifaceted and complex, touching upon ethics, privacy, technical integration, legal frameworks, environmental impact, regulatory compliance, and workforce readiness. While the potential benefits are significant, navigating these challenges requires a comprehensive and strategic approach. Companies must be proactive in addressing these issues by investing in compliance, sustainability, and workforce development, as well as by establishing ethical guidelines and frameworks that govern the use of AI and immersive technologies.

As these technologies continue to evolve and become more integral to business operations, businesses that effectively manage these challenges will not only gain a competitive edge but also set the standard for responsible and sustainable innovation in the digital economy.

The Future of Generative AI and Spatial Computing in Business

The trajectory of Generative AI and Spatial Computing suggests a future where businesses become more adaptive, immersive, and efficient. The use of digital twins is expected to grow, not just in smart cities but across various sectors like manufacturing and healthcare. Businesses will increasingly simulate entire processes or products in digital environments before physical implementation, reducing waste and improving quality.

As technology matures, the cost and complexity of integrating Generative AI and Spatial Computing solutions will likely decrease, allowing more businesses, including SMEs, to access these innovations.

AI's ability to understand and anticipate customer needs will enable companies to deliver even more personalized and immersive experiences. This trend will be critical for businesses focusing on customer retention and engagement.

The evolution of Generative AI-powered robotics and automation will be intertwined with Spatial Computing, enabling businesses to automate complex physical tasks. For example, AI-generated environments could be used to train autonomous robots in simulated conditions before deployment in real-world operations.

The combination of Generative AI and Spatial Computing is not merely a technological innovation but a transformative force that reshapes business strategies and operations. Companies that invest in and adapt to these technologies will gain competitive advantages, create new market opportunities, and achieve higher efficiency. However, navigating the challenges related to cost, ethics, and technical integration will be crucial for long-term success. As the technologies continue to mature, their integration will become essential for businesses looking to stay ahead in a rapidly evolving digital economy.

In the next chapter, we discuss the considerable impact of the Apple Vision Pro on both consumers and businesses alike.

13

The Impact of the Apple Vision Pro

The Apple Vision Pro, unveiled in 2023, represents the Apple's groundbreaking entry into the rapidly evolving world of MR and Spatial Computing. Positioned as a cutting-edge wearable device, the Vision Pro combines AR and VR to deliver a seamless Spatial Computing experience. Unlike traditional headsets, this device seeks to blend physical and digital worlds, enabling users to interact with applications and environments in an entirely immersive manner. Apple's announcement of the Vision Pro signals its commitment to revolutionizing how people engage with technology in their daily lives.

Equipped with visionOS, Apple's first operating system designed specifically for Spatial Computing, the Vision Pro offers a robust suite of capabilities. These include gesture control, eye-tracking, and voice input, redefining user interaction and interface design. With its advanced micro-OLED displays, users experience unparalleled visual fidelity, making content appear lifelike and immersive. The device's outward-facing cameras allow for passthrough functionality, enabling users to remain connected to their physical surroundings while navigating digital spaces. This blending of realities ensures the Vision Pro isn't just another VR headset but a device poised to establish a new standard for immersive technology.

During its launch, Apple positioned the Vision Pro as a versatile tool catering to professionals, creators, and everyday users alike. From transforming workspaces with virtual monitors to providing immersive entertainment like 3D movies and interactive gaming, the Vision Pro is more than a device—it's Apple's vision for the future of computing. Despite its steep initial price point, the Vision Pro has garnered significant attention for its potential to shape multiple industries, signaling Apple's bold step into a market previously dominated by Meta, Microsoft, and other players.

In this chapter, we discuss technological features, Apple Vision Pro applications across industries, content creation and development, economic and market impact, challenges and criticisms, the future of Spatial Computing, case studies and early use cases, and Apple Vision Pro's future.

Technological Features

VisionOS and User Interface Innovations

At the heart of the Vision Pro is visionOS, which offers an intuitive interface that adapts to user behavior, enabling interactions that feel natural and effortless. Unlike traditional 2D operating systems, visionOS operates within a 3D environment, allowing users to interact with applications and digital content as if they were physical objects in their surroundings due to full-color, high-resolution passthrough.

VisionOS leverages spatial navigation, allowing users to position, resize, and manipulate apps in their field of view. This makes multitasking in immersive environments seamless and highly efficient.

The operating system supports native applications tailored for Spatial Computing while maintaining compatibility with existing Apple apps. Popular tools like Safari, FaceTime, and Photos have been reimagined to function dynamically within a mixed-reality space.

Eye-Tracking, Gesture Control, and Passthrough Capabilities

Apple has revolutionized user interaction with the Vision Pro's combination of eye-tracking, gesture control, and voice input, eliminating the need for external controllers.

Advanced eye-tracking sensors follow the user's gaze with pinpoint accuracy, enabling natural navigation and interaction. For instance, users can select items by simply looking at them, a functionality that feels both intuitive and futuristic.

The Vision Pro uses cameras to detect hand gestures, allowing users to scroll, zoom, and interact with objects in real time. The fluidity of these interactions eliminates friction, ensuring a highly responsive experience.

Unlike many VR headsets that isolate users from their surroundings, the Vision Pro features passthrough capabilities, enabling users to maintain awareness of their physical environment. High-resolution external cameras stream the real world into the user's view, blending AR and VR seamlessly.

Micro-OLED Displays and Hardware Advancements

The Vision Pro's hardware is engineered for unparalleled performance and comfort. Its micro-OLED display delivers stunning visual fidelity, with over 23 million pixels spread across the two lenses. This ensures that digital content appears crisp, vibrant, and lifelike.

The micro-OLED technology allows for ultra-high resolution, offering 4K clarity for each eye. This creates an immersive experience ideal for watching movies, playing games, or working in virtual environments.

It is powered by Apple's M2 chip and a custom R1 chip, enabling low-latency processing for real-time interactions. This dual-chip architecture ensures smooth performance, even when running demanding applications.

Apple has prioritized user comfort, incorporating lightweight materials and an adjustable headband for prolonged use. The device's modular design allows for customization, ensuring it fits a variety of users comfortably.

Apple's Vision for Immersive Technology

The Vision Pro represents Apple's bold vision for the future of computing. By introducing Spatial Computing, Apple aims to blur the lines between the physical and digital worlds, creating an ecosystem where technology becomes a natural extension of human interaction.

Reimagining Computing Paradigms

Apple envisions a future where computing transcends the limitations of screens and devices. With the Vision Pro, interactions occur directly within the user's space, making technology less intrusive and more integrated into everyday life. This paradigm shift aligns with Apple's long-term goal of creating devices that enhance human productivity and creativity without disrupting natural behavior.

Developer Empowerment

To realize this vision, Apple has invested heavily in tools for developers. VisionOS **software development kits (SDKs)** enable creators to build immersive applications tailored for Spatial Computing. By fostering a robust developer ecosystem, Apple ensures a steady stream of innovative applications that take full advantage of the Vision Pro's capabilities. Recently, Apple has started to allow the WebXR API to run on Apple Vision Pro's Safari browsers, allowing non-Apple developers to build MR content for the Apple Vision Pro.

Ecosystem Integration

Apple's focus on ecosystem integration ensures that the Vision Pro is not a standalone device but part of a larger interconnected network. Features like iCloud synchronization, seamless pairing with Macs and iPads, and compatibility with AirPods create a cohesive **user experience (UX)**. This integration reinforces Apple's commitment to creating technology that works harmoniously across all its products.

Inspiring a New Era of Content Creation

The Vision Pro opens new doors for creators, offering tools to design 3D content, immersive films, and interactive experiences. Apple's vision extends to empowering artists, filmmakers, and designers to explore Spatial Computing as a new medium for creativity.

The Apple Vision Pro stands as a testament to Apple's innovative approach to Spatial Computing. By combining cutting-edge hardware, intuitive software, and a focus on seamless integration, Apple has redefined what an MR device can achieve. Its advanced features, differentiation from competitors, and alignment with Apple's broader vision for technology highlight the Vision Pro's potential to shape the future of immersive experiences. As the Vision Pro sets new standards, it serves as a catalyst for the growth of Spatial Computing, paving the way for transformative applications across industries and everyday life.

We will now discuss how the Apple Vision Pro deals with interaction.

Redefining Interaction in Spatial Computing

The Apple Vision Pro represents a paradigm shift in how humans interact with digital technology. Moving beyond traditional interfaces like screens and keyboards, the Vision Pro introduces innovative interaction mechanisms that blur the boundaries between the physical and digital worlds. By leveraging **natural user interfaces (NUIs)**, immersive spatial environments, and accessible UX design, Apple has set the stage for Spatial Computing to become a cornerstone of future technology.

Natural User Interfaces

NUIs aim to make interaction with digital systems intuitive by aligning with human behaviors such as gestures, voice commands, and eye movements. The Vision Pro's suite of NUI technologies, including eye-tracking, gesture recognition, and voice input, introduces a seamless, immersive way to engage with digital content.

Eye-Tracking: Precision and Context Awareness

Eye-tracking is one of the Vision Pro's most innovative features, enabling users to navigate and interact with applications by simply looking at objects on the screen. Cameras within the headset monitor the user's gaze, translating their focus into actionable commands.

Users can select menu options or open applications by gazing at icons, eliminating the need for physical controllers or touchscreens.

Eye-tracking allows the Vision Pro to understand where users are looking, providing contextual prompts or adjusting content to match their focus. For instance, during a presentation, the device can prioritize content in the user's line of sight while dimming peripheral elements.

For individuals with physical disabilities, eye-tracking offers an inclusive method of interacting with technology, breaking down traditional barriers.

Gesture Recognition: Intuitive Interaction

The Vision Pro uses an array of external cameras to detect hand movements and gestures, enabling a highly responsive interaction system. Unlike traditional AR/VR systems that rely on handheld controllers, the Vision Pro lets users interact with digital content naturally using their hands, performing much better than other XR devices that recognize hand gestures.

Simple gestures, such as pinching to zoom or swiping to scroll, mimic familiar physical interactions, making the technology more intuitive for new users.

Gesture controls allow users to manipulate multiple applications simultaneously, such as resizing a window with one hand while scrolling through content with the other.

Gesture recognition is not confined to a limited field, offering users the flexibility to interact with the device from various positions.

Voice Input: Conversational Control

Voice input further expands the Vision Pro's NUI capabilities, allowing users to issue commands, search for information, or dictate text without lifting a finger.

The device's advanced NLP algorithms interpret spoken commands with high accuracy, creating a conversational interaction experience.

Voice commands can be combined with eye-tracking and gestures for more complex interactions. For example, a user might look at an object, point to it, and issue a verbal command to manipulate it.

Voice input enhances productivity in scenarios where hands-free operation is essential, such as during cooking, workouts, or industrial tasks.

The integration of eye-tracking, gestures, and voice input not only redefines user interaction but also lays the groundwork for a future where technology adapts seamlessly to human behavior.

Transition from Screens to Spaces

The Vision Pro represents a significant leap from screen-based interactions to immersive spatial interfaces, redefining how digital content is experienced. This transition challenges traditional paradigms by embedding technology into the user's environment, creating a more natural and engaging computing experience.

From Flat Screens to 3D Interfaces

Conventional screens, whether on smartphones, laptops, or TVs, restrict interaction to two dimensions. The Vision Pro breaks free of these limitations by rendering digital content in 3D space.

Applications and media are no longer confined to rectangular displays. Instead, they exist as part of the user's surroundings, enhancing immersion and engagement.

Users can position virtual monitors, calendars, or tools anywhere in their field of view, tailoring their environment to suit their preferences and workflows.

Unlike static screens, spatial interfaces allow users to walk around, inspect, and interact with digital objects as if they were physical and they remain persistent (for example, if you moved to another room and then returned to the first room, the spatial interface would be in the exact same place where you left it).

Passthrough and MR

The Vision Pro's passthrough capabilities blur the boundaries between physical and digital spaces, enabling users to interact with both simultaneously.

By overlaying digital content onto the real world, users can interact with virtual tools while staying connected to their environment. For instance, architects can visualize 3D building models directly on construction sites.

Unlike fully immersive VR systems, the Vision Pro keeps users aware of their surroundings, ensuring safety and reducing the risk of disorientation.

Collaboration in Shared Spaces

Spatial Computing also transforms collaborative experiences, enabling multiple users to interact in shared virtual environments.

Teams can meet in immersive spaces where 3D models, presentations, and documents can be manipulated collaboratively.

The Vision Pro's ability to integrate with other Apple devices ensures a seamless transition between physical and virtual workflows, enhancing productivity and creativity.

By transitioning from screens to spaces, the Vision Pro not only changes how users engage with technology but also reimagines the environments in which they work, play, and collaborate.

UX Design

The success of any technological innovation hinges on its UX, and the Vision Pro sets a new standard for Spatial Computing UX. Apple's commitment to intuitive design and accessibility ensures the device is not just powerful but also user-friendly.

Intuitive Design Principles

The Vision Pro's UX is grounded in simplicity and intuitiveness, making advanced Spatial Computing accessible to a broad audience.

The use of natural gestures, eye-tracking, and voice commands minimizes the need for extensive tutorials, allowing users to start interacting with the device immediately.

Subtle animations and visual feedback guide users through interactions, ensuring clarity and reducing the likelihood of errors.

The Vision Pro adapts to individual user preferences, learning from behavior to optimize layout, shortcuts, and content presentation.

Accessibility Features

Apple has a long history of prioritizing accessibility, and the Vision Pro continues this tradition by ensuring that Spatial Computing is inclusive.

Users with mobility impairments can rely solely on voice or eye-tracking for interaction, bypassing the need for hand gestures.

Adjustable font sizes, color contrasts, and interface layouts accommodate users with visual impairments or other disabilities.

Built-in accessibility tools provide step-by-step guidance, ensuring that all users can navigate and benefit from the Vision Pro's features.

Reducing Cognitive Load

Spatial Computing introduces complexities that could overwhelm users, but the Vision Pro UX is designed to reduce cognitive load.

Applications dynamically adjust their complexity based on user behavior, displaying only relevant information and tools.

The Vision Pro allows users to manage multiple tasks without clutter or confusion, streamlining workflows and boosting productivity.

Advancing UX Design in Spatial Computing

The Vision Pro's UX design is not just about improving the UX—it's about setting a benchmark for the Spatial Computing industry.

Apple's design principles encourage developers to prioritize intuitive and accessible interfaces for their visionOS applications.

As competitors look to Apple's innovations for inspiration, the Vision Pro is likely to influence the design of future Spatial Computing devices.

The Apple Vision Pro redefines interaction in Spatial Computing by integrating NUIs, shifting from screens to immersive spaces, and elevating UX design. Eye-tracking, gestures, and voice input create intuitive and seamless interactions, while spatial interfaces revolutionize how users engage with digital content in their environments. By prioritizing accessibility and reducing cognitive load, Apple ensures that the Vision Pro is not only innovative but also inclusive. This reimagining of human-computer interaction paves the way for Spatial Computing to become a transformative force in everyday life and professional workflows.

Applications of the Vision Pro across Industries

The Apple Vision Pro is a versatile tool poised to transform multiple industries through its cutting-edge technology and Spatial Computing capabilities. With its ability to blend AR and VR, the Vision Pro introduces innovative ways for businesses, professionals, and consumers to engage with digital content. By leveraging immersive experiences, real-time interactivity, and seamless integration into workflows, the Vision Pro unlocks new possibilities in entertainment, productivity, healthcare, education, and retail.

Entertainment and Media

Immersive Streaming and 3D Cinema Experiences

The Vision Pro is set to revolutionize entertainment by bringing immersive streaming and 3D cinema experiences directly to users. With its high-resolution micro-OLED displays and spatial audio technology, the Vision Pro delivers a cinematic experience that rivals traditional theaters.

Users can watch movies and TV shows in 3D, with virtual screens that adjust in size and location to fit their preferences. This makes every viewing session feel like a private cinema experience.

Platforms can offer immersive content, such as live concerts or sports events, where viewers feel as though they are part of the action. The Vision Pro's ability to render real-world environments adds an extra layer of realism to these experiences.

Directors and creators can design interactive films that allow viewers to engage with the storyline, choosing perspectives or interacting with characters in a virtual environment.

Gaming in Augmented and Mixed Reality

The Vision Pro's AR and mixed-reality capabilities redefine gaming, providing players with immersive environments that integrate seamlessly with their physical surroundings.

Players can interact with digital objects in their real-world environment, creating unique gaming experiences. For example, users can engage in a tabletop game where virtual characters and objects appear on their actual table.

The Vision Pro enables multiplayer gaming experiences where players from different locations can interact in shared virtual spaces. This fosters collaboration and competition on an entirely new level.

Combined with peripherals like haptic gloves or controllers, the Vision Pro deepens immersion by providing tactile feedback during gameplay.

Productivity and Collaboration

Virtual Workspaces and Remote Collaboration Tools

The Vision Pro transforms traditional workflows by introducing virtual workspaces that adapt to the user's needs. Professionals can set up an unlimited number of virtual screens and tools, enhancing productivity and multitasking.

Users can create virtual desktops that follow them wherever they go, enabling seamless transitions between physical and digital environments. This eliminates the need for multiple physical monitors.

Teams can meet in virtual spaces equipped with interactive 3D tools, such as whiteboards or prototypes. Vision Pro's spatial audio and high-resolution visuals ensure that remote interactions feel as natural as in-person meetings, utilizing "Personas," which are AI representations of what a person looks like without their headset on.

Users can share files, presentations, or models in a virtual space, enabling real-time collaboration and feedback.

Enhanced Presentations and Prototyping in 3D Spaces

The Vision Pro empowers professionals to create and present ideas in ways that were previously impossible.

Presenters can use it to display interactive 3D models, such as architectural designs or product prototypes, allowing audiences to explore them from every angle.

Designers and engineers can create and manipulate 3D prototypes in virtual spaces, testing concepts without the need for physical materials. This reduces costs and accelerates the design process.

Businesses can use the Vision Pro to deliver compelling presentations to clients, creating a memorable and engaging experience that sets them apart from competitors.

Healthcare

Training and Simulations for Medical Professionals

The Vision Pro introduces groundbreaking applications in healthcare, particularly in training and simulations.

Medical professionals can use it to practice complex procedures in a risk-free virtual environment. Realistic simulations allow for hands-on learning without the need for a physical patient.

First responders can participate in virtual drills that simulate real-life emergencies, preparing them to act decisively under pressure.

Students and professionals can explore detailed 3D models of the human body, enhancing their understanding of anatomy and physiology.

Patient Care and Rehabilitation through Immersive Experiences

Beyond training, the Vision Pro can be used to improve patient care and rehabilitation outcomes.

Patients undergoing physical therapy can use it to engage in guided exercises, with real-time feedback provided by the device. Virtual environments can make therapy sessions more engaging and enjoyable.

Immersive environments can be used to treat conditions such as anxiety or PTSD, providing calming virtual experiences or controlled exposure therapy.

The Vision Pro allows doctors to consult with patients in virtual spaces, enabling more interactive and personalized telemedicine experiences.

Education

Transforming Classrooms and Virtual Learning with Interactive 3D Content

The Vision Pro has the potential to redefine education by making learning more interactive and engaging.

Students and teachers can connect in immersive virtual spaces, where lessons come to life through 3D visualizations and simulations. For example, a history class could virtually explore ancient civilizations.

Subjects like chemistry or engineering can be taught using virtual labs, allowing students to experiment safely and without the need for expensive equipment.

Vision Pro enables students from around the world to access high-quality education, bridging geographical and socioeconomic divides.

Democratizing Access to Immersive Educational Tools

By leveraging Apple's ecosystem, the Vision Pro can make advanced educational tools accessible to a wider audience.

Students can practice languages by interacting with virtual tutors or exploring simulated environments in the target language's culture.

Fields such as aviation or space exploration can use it to provide specialized training to students and professionals in cost-effective ways.

The Vision Pro's interactive capabilities make STEM subjects more engaging, inspiring students to pursue careers in science and technology.

Retail and E-Commerce

AR-Powered Shopping and Virtual Product Experiences

The Vision Pro introduces innovative ways for customers to shop and engage with products.

Shoppers can use it to visualize how furniture, clothing, or other products would look in their homes or on their bodies. This reduces uncertainty and improves purchase confidence.

Retailers can create virtual showrooms where customers can explore products in immersive 3D environments. This enhances the shopping experience without requiring physical inventory.

Brands can use the Vision Pro to offer virtual product demonstrations, allowing customers to see and interact with products before making a purchase.

Enhancing Customer Engagement through Spatial Computing

The Vision Pro redefines how brands interact with their customers, creating more meaningful and engaging experiences.

Retailers can use it to create tailored experiences for customers, such as virtual styling sessions or customized product recommendations.

Brands can gamify loyalty programs by integrating AR elements, encouraging repeat engagement and fostering brand loyalty.

For physical stores, it can provide interactive in-store experiences, such as guided tours or AR-based promotions.

The Apple Vision Pro is poised to make a significant impact across industries, offering new ways to engage, innovate, and create value. From revolutionizing entertainment and gaming to enhancing productivity, education, and healthcare, the Vision Pro is a tool for transformation. Its ability to merge the physical and digital worlds opens up endless possibilities, making it a pivotal player in the future of Spatial Computing. As businesses and consumers embrace this technology, the Vision Pro is set to redefine how industries operate and evolve in the years to come.

Transforming Content Creation and Development with the Apple Vision Pro

The Apple Vision Pro is revolutionizing content creation by introducing tools and capabilities that empower filmmakers, designers, developers, and other creative professionals. This section explores the creative possibilities enabled by the Vision Pro, its impact on spatial app development, and its role in establishing a thriving 3D content ecosystem.

Creative Possibilities: Redefining Opportunities for Filmmakers, Designers, and Developers

The Vision Pro unlocks new creative horizons by enabling professionals to reimagine storytelling, design, and digital interaction in spatial environments. From filmmakers creating immersive narratives to designers crafting interactive 3D models, its capabilities extend the boundaries of traditional content creation.

For Filmmakers: Immersive Storytelling

Filmmakers can leverage the Vision Pro to craft stories that go beyond flat screens, immersing viewers in narratives where they can interact with the environment or choose their perspectives.

Filmmakers can produce content that places viewers at the center of the action, allowing them to explore different angles and scenes as the story unfolds. This offers a level of engagement that traditional cinema cannot achieve.

The Vision Pro allows filmmakers to incorporate interactivity into their storytelling. For example, users can influence the storyline by making choices or interacting with characters, creating personalized experiences.

Directors and cinematographers can use it to visualize and pre-plan shots in a virtual environment before actual filming, reducing production costs and improving efficiency.

For Designers: Enhanced 3D Modeling and Visualization

The Vision Pro transforms how designers approach 3D modeling, product design, and visualization by offering tools that bring their creations to life in spatial environments.

Designers can build and manipulate 3D models in real time, walking around their creations and examining them from every angle. This capability is invaluable in fields like architecture, automotive design, and product development.

Teams can collaborate on designs in shared virtual spaces, enabling real-time feedback and adjustments. For example, an automotive design team can work together on a virtual car prototype, tweaking details while visualizing the final product.

Designers can present their creations to clients or stakeholders in immersive 3D spaces, offering a more compelling way to communicate ideas and concepts.

For Developers: Pioneering Interactive Applications

Developers are at the forefront of leveraging the Vision Pro's capabilities to create spatial applications that redefine user interaction.

Developers can build apps that integrate digital content into real-world environments, such as AR-based navigation tools or interactive gaming experiences.

The Vision Pro enables developers to create applications that seamlessly integrate with other Apple devices, ensuring a cohesive UX across the ecosystem.

By combining the Vision Pro's capabilities with machine learning, developers can create intelligent applications that adapt to user behavior, such as personalized educational tools or interactive fitness apps.

The creative potential of the Vision Pro spans industries and disciplines, making it a transformative tool for professionals seeking to push the boundaries of their craft.

Spatial App Development: Empowering Developers with VisionOS and Apple's Tools

Apple has provided developers with a suite of tools and resources to harness the Vision Pro's capabilities, ensuring the creation of innovative and impactful spatial applications. VisionOS serves as the foundation for spatial app development, enabling developers to design experiences that fully utilize the device's features.

Developer Tools and SDKs

Apple has introduced a range of tools and SDKs to empower developers in creating spatial applications.

This tool enables developers to design and prototype AR and VR experiences with ease. From placing 3D objects in virtual environments to testing interactions, Reality Composer Pro simplifies the creation process.

Apple's ARKit framework and Metal graphics API provide developers with the technical capabilities to build high-performance, visually stunning applications for the Vision Pro.

Apple's developer tools include capabilities for creating multi-user applications, allowing developers to design collaborative experiences in shared virtual spaces.

Opportunities for Developers

The Vision Pro opens up new avenues for developers to create applications across various industries, from education and healthcare to entertainment and productivity.

Businesses can collaborate with developers to create bespoke applications tailored to their specific needs, such as virtual training modules or interactive product demonstrations.

Developers can leverage the Vision Pro's app ecosystem to market and sell their applications, tapping into a growing user base interested in Spatial Computing. By contributing innovative applications to its app store, developers play a crucial role in expanding the device's functionality and appeal.

Through visionOS and Apple's developer tools, the Vision Pro empowers developers to create applications that redefine user interaction and unlock new possibilities in Spatial Computing.

3D Content Ecosystem: Promoting Mass Adoption of 3D Content Creation

The Vision Pro is not just a device for consuming 3D content—it is a catalyst for building a thriving ecosystem that encourages the creation and dissemination of immersive experiences. By promoting the adoption of 3D content creation, Apple aims to make Spatial Computing a mainstream technology.

Lowering Barriers to Entry

The Vision Pro simplifies the process of creating 3D content, making it accessible to a wider audience.

Tools like Reality Composer Pro empower creators with limited technical expertise to design 3D content, democratizing access to Spatial Computing.

Creators can use familiar Apple devices, such as iPads or Macs, to design 3D content that integrates seamlessly with the Vision Pro.

Apple's resources include pre-built templates and step-by-step tutorials, reducing the learning curve for new creators.

Fostering Collaboration

The Vision Pro's ecosystem encourages collaboration between creators, developers, and businesses.

Teams can work together on 3D projects in virtual environments, combining their expertise to produce high-quality content.

Apple's ecosystem includes a marketplace where creators can share or sell 3D models, textures, and other assets, fostering a community of innovations.

Driving Consumer Adoption

By promoting the creation of high-quality 3D content, the Vision Pro paves the way for greater consumer adoption of Spatial Computing.

A growing library of 3D content ensures that users have access to a wide range of applications and experiences, from entertainment and education to productivity and healthcare.

As more creators contribute to the ecosystem, the Vision Pro becomes increasingly appealing to mainstream audiences, driving its adoption across demographics.

The availability of diverse content inspires developers and creators to experiment with new ideas, pushing the boundaries of what Spatial Computing can achieve.

Through its support for 3D content creation, the Vision Pro establishes a self-sustaining ecosystem that drives innovation, collaboration, and mass adoption of Spatial Computing.

Next, we tackle the significance of Apple Vision Pro on economics and markets.

Economic and Market Impact of the Apple Vision Pro

The Apple Vision Pro marks a pivotal moment in the development and commercialization of Spatial Computing technologies. As a premium MR device, it not only sets a new benchmark in terms of capabilities but also has the potential to influence consumer adoption, reshape competition, and create novel market opportunities.

Consumer Adoption: Driving Spatial Computing into the Mainstream

The Apple Vision Pro enters the market as a high-end device with a price tag that reflects its advanced technology and premium features. Despite its steep cost, it has the potential to drive widespread consumer adoption of Spatial Computing by leveraging Apple's brand loyalty, ecosystem integration, and focus on UX.

Appealing to Early Adopters and High-End Consumers

The Vision Pro is currently targeted at early adopters and professionals willing to invest in cutting-edge technology. This group has historically driven the success of Apple's innovations, such as the iPhone and Apple Watch, paving the way for wider adoption.

High-end consumers are likely to embrace the Vision Pro for its groundbreaking capabilities in productivity, entertainment, and creativity. These early users will serve as advocates, showcasing its potential to mainstream audiences.

As one of the most trusted and influential technology brands, Apple has a proven track record of transforming niche markets into mainstream phenomena. The Vision Pro benefits from Apple's reputation for high-quality products and its ability to shape consumer expectations.

Lowering Barriers to Entry over Time

While the Vision Pro's initial price point may limit its audience, Apple's long-term strategy likely includes making the technology more accessible.

As manufacturing processes become more efficient and components like micro-OLED displays and processors are produced at scale, Apple could introduce more affordable versions of the Vision Pro, similar to how the iPhone SE broadened access to the iPhone ecosystem.

As developers create applications tailored to various demographics—such as students, educators, and small businesses—the Vision Pro's value proposition will appeal to a broader audience.

Mainstream Appeal through Ecosystem Integration

The Vision Pro's seamless integration with Apple's ecosystem further enhances its appeal to consumers.

Existing Apple users will appreciate the Vision Pro's ability to work seamlessly with their iPhones, iPads, and Macs, creating a cohesive UX.

Features such as iCloud integration, shared content across devices, and compatibility with Apple services (e.g., Apple TV and Apple Fitness+) increase the Vision Pro's value, encouraging adoption within Apple's loyal customer base.

By targeting early adopters, leveraging its ecosystem, and potentially broadening its reach with future iterations, the Vision Pro has the potential to accelerate the adoption of Spatial Computing technologies.

Impact on Competitors: Shaping the Spatial Computing Market

Apple's entry into the Spatial Computing market with the Vision Pro is set to disrupt the competitive landscape, forcing existing players to rethink their strategies. Companies like Meta, Microsoft, and HTC, which have led the AR/VR market, now face a formidable competitor with unmatched resources, innovation capabilities, and market influence.

Setting a New Benchmark for Quality

The Vision Pro redefines the standard for MR devices, particularly in terms of display quality, interaction mechanisms, and UX.

By offering features such as micro-OLED displays with ultra-high resolution, advanced eye-tracking, and seamless ecosystem integration, Apple has positioned the Vision Pro as the gold standard for Spatial Computing.

Competitors like Meta, which has focused on affordability with devices like the Quest series, may need to innovate in areas like visual fidelity and interaction to remain competitive.

Driving Innovation across the Industry

Apple's entry into Spatial Computing is likely to spark a wave of innovation across the industry as competitors seek to match or surpass the Vision Pro's capabilities.

The Vision Pro's focus on intuitive interactions and polished design will push competitors to prioritize UX in their devices, moving beyond hardware specifications alone.

The Vision Pro's success may encourage greater investment in AR and VR technologies, leading to advancements in areas like battery life, processing power, and content creation.

Expanding the Market

Apple's presence in the Spatial Computing market could expand its overall size and appeal, benefiting all players.

Apple's marketing power will bring mainstream attention to Spatial Computing, increasing consumer curiosity and interest in the category as a whole.

Competitors and partners may benefit from collaborating with Apple to develop applications and standards that accelerate adoption across the industry.

While the Vision Pro presents challenges for existing competitors, it also has the potential to elevate the entire Spatial Computing market by driving innovation and increasing consumer awareness.

Market Opportunities: Unlocking New Business Models and Ecosystems

The Vision Pro's launch represents more than a product release—it signals the emergence of new business models, application ecosystems, and industries driven by Spatial Computing. Apple's emphasis on creating a robust ecosystem around the Vision Pro is likely to inspire innovative solutions and revenue streams.

Monetizing Spatial Applications

The Vision Pro opens up opportunities for developers and businesses to create spatial applications tailored to a wide range of use cases.

Apple's dedicated app store for Vision Pro applications will serve as a hub for developers to market and sell their creations, generating new revenue streams.

Businesses can offer subscription-based services, such as virtual training modules, immersive fitness programs, or AR-assisted design tools, leveraging the Vision Pro's capabilities.

Developers can incorporate in-app purchases for premium features, 3D assets, or content packs, further diversifying their revenue opportunities.

Fostering Collaboration and Innovation

The Vision Pro's ecosystem encourages collaboration between developers, businesses, and creatives, resulting in innovative solutions.

Apple's developer tools, such as Reality Composer Pro and visionOS SDKs, provide a foundation for collaborative innovation in Spatial Computing.

Apple's partnerships with businesses in healthcare, education, and entertainment will drive the development of tailored applications and solutions.

Accelerating Content Creation

The Vision Pro's focus on immersive experiences will catalyze the creation of 3D content for both professional and consumer markets.

The rise of spatial applications will increase demand for tools and platforms that simplify 3D content creation, opening up opportunities for software developers and hardware manufacturers.

The Vision Pro could usher in a new era of mainstream 3D media, from immersive films to interactive gaming experiences.

The Apple Vision Pro's economic and market impact extends far beyond its initial launch. By driving consumer adoption, setting new standards for quality and innovation, and unlocking new business models and ecosystems, the Vision Pro has the potential to reshape the Spatial Computing landscape. Apple's entry into this market not only challenges competitors but also expands opportunities for developers, businesses, and industries to embrace immersive technology. As the Vision Pro ecosystem evolves, it will play a pivotal role in defining the future of computing and creating value across the global economy.

Challenges and Criticisms of the Apple Vision Pro

While the Apple Vision Pro introduces groundbreaking advancements in Spatial Computing, it is not without its challenges and criticisms. These issues, ranging from technical limitations to ethical concerns, highlight the complexities of adopting such a transformative technology. Addressing these challenges will be critical for the Vision Pro to achieve widespread adoption and long-term success.

Technical Limitations

The Vision Pro represents a significant leap forward in hardware and software design, but it also faces several technical limitations. Issues such as battery life, device weight, and processing power underscore the need for further innovation to refine its performance and usability.

Battery Life

One of the most immediate technical challenges for the Vision Pro is its battery life, which is reported to last approximately two hours on a full charge when using the external battery pack.

A two-hour battery life restricts the Vision Pro's utility, particularly for use cases like extended work sessions, long-form content consumption, or immersive gaming.

The need for an external battery pack adds to the device's complexity, detracting from the sleek and self-contained design Apple is known for.

Apple could explore innovations in battery technology, such as solid-state batteries or wireless charging capabilities, to extend usage time and improve convenience.

Weight and Comfort

The Vision Pro's sophisticated hardware comes at the cost of added weight, which may affect user comfort during extended use.

The weight of the device, combined with its head-mounted design, may cause physical discomfort for users, particularly when worn for long periods.

Apple must strike a balance between incorporating advanced features and maintaining a lightweight, ergonomic design.

Materials such as carbon fiber or magnesium alloys, as well as optimized weight distribution, could help make the Vision Pro more comfortable for everyday use.

Processing Power and Heat Management

The Vision Pro's high-performance capabilities place significant demands on its processing power, leading to concerns about heat generation and system performance. Sustained use may result in overheating, potentially affecting performance and user comfort.

While the M2 and R1 chips are powerful, future iterations may require even more advanced processors to handle complex Spatial Computing tasks seamlessly.

As developers create more demanding applications, the Vision Pro's hardware may struggle to keep pace without further advancements in chip technology.

Adoption Barriers

The Vision Pro's potential to transform Spatial Computing is tempered by barriers to adoption, particularly its high cost and limited accessibility for mainstream consumers.

Cost Concerns

Priced at \$3,499, the Vision Pro is positioned as a premium product, placing it out of reach for many potential users.

The high price point restricts adoption to a niche audience of early adopters, professionals, and high-income consumers.

Competing devices, such as Meta's Quest series, offer more affordable options, making it challenging for the Vision Pro to compete in terms of price.

To expand its user base, Apple will need to introduce lower-cost variants or financing options to make the Vision Pro more accessible.

Accessibility for Diverse Demographics

Beyond cost, accessibility concerns extend to how the Vision Pro accommodates users with varying needs and abilities.

Limited availability in certain regions may hinder its global adoption, particularly in emerging markets where disposable income is lower.

While Apple is known for its accessibility features, ensuring the Vision Pro is usable by individuals with disabilities will be essential for widespread adoption.

The high cost and technical requirements could exacerbate the digital divide, leaving economically disadvantaged communities behind.

Privacy and Ethical Considerations

As an advanced Spatial Computing device, the Vision Pro collects vast amounts of user data, raising significant privacy and ethical concerns. These issues must be addressed to build trust and ensure responsible use of the technology.

Risks Associated with Eye-Tracking and Data Collection

The Vision Pro's eye-tracking capabilities introduce new opportunities for personalization but also raise concerns about data security and user privacy.

Eye-tracking data can reveal highly personal information, such as emotional states, interests, and attention patterns, making it a target for misuse or unauthorized access. The collection and storage of such sensitive data could be exploited for surveillance purposes, either by third parties or malicious actors.

Apple must ensure that its data collection practices comply with privacy regulations like GDPR and CCPA, while also being transparent about how user data is stored and used.

Ethical Implications of Immersive Technology on User Well-Being

The immersive nature of the Vision Pro raises questions about its impact on mental and physical health, as well as its broader societal implications.

Immersive technologies can be highly engaging, potentially leading to excessive use and negative effects on mental health, such as reduced social interaction or escapism.

Extended use may contribute to issues like eye strain, neck discomfort, or posture problems, particularly if users do not take regular breaks.

Developers and content creators have a responsibility to ensure that applications designed for the Vision Pro promote positive and ethical experiences, avoiding exploitative or harmful content.

Balancing Innovation and Responsibility

To address these ethical concerns, Apple must take a proactive approach to ensuring it is used responsibly.

Built-in features such as usage reminders, parental controls, and health monitoring can help mitigate the risks of overuse or misuse.

Apple should educate users about the potential risks of immersive technology and provide clear guidelines on how to use it safely.

By embedding ethical considerations into the design and development of the Vision Pro, Apple can set a precedent for responsible innovation in Spatial Computing.

Addressing these challenges will require a combination of technological advancements, strategic pricing, and a commitment to transparency and ethical design. By tackling these issues head-on, Apple has the opportunity to not only refine the Vision Pro but also set a standard for the future of Spatial Computing. As the Vision Pro evolves, its success will depend on Apple's ability to balance innovation with responsibility, ensuring that this groundbreaking device delivers on its promise to transform the way we interact with technology.

We will now discuss how the Apple Vision Pro impacts the future of Spatial Computing.

Vision Pro and the Future of Spatial Computing

The Apple Vision Pro represents a groundbreaking step toward the future of Spatial Computing, combining cutting-edge AR, VR, and MR capabilities into a single device. With its advanced hardware, seamless software integration, and potential to interface with emerging technologies, it is not only a product but a catalyst for the development of the Spatial Computing ecosystem. This section explores how it advances MR, integrates with emerging technologies, and contributes to the broader growth of Spatial Computing, including its foundational role in building the “spatial web.”

Advancing Mixed Reality

The Vision Pro's design and capabilities demonstrate a leap forward in MR, blending digital and physical environments in ways that were previously unattainable. It sets a new standard for MR devices by offering superior interaction mechanisms, immersive visuals, and versatile applications that pave the way for future innovation.

Blurring the Lines between AR and VR

The Vision Pro's ability to fluidly transition between AR and VR is a hallmark of its advanced MR capabilities.

Users can switch effortlessly from viewing their physical surroundings enhanced with digital overlays (AR) to fully immersive virtual environments (VR), allowing for flexible applications across entertainment, productivity, and communication.

External cameras enable passthrough functionality, providing users with a high-fidelity view of their physical environment. This capability is essential for creating immersive mixed-reality experiences where digital content interacts meaningfully with the real world.

Setting a New Benchmark for Hardware

Spatial audio technology ensures that sound appears to originate from specific locations within the user's environment, enhancing immersion and realism.

The device's eye-tracking and gesture recognition enable intuitive and fluid interactions with digital objects, eliminating the need for physical controllers.

Inspiring Future MR Devices

The Vision Pro's success is likely to influence the development of more sophisticated MR devices.

By introducing premium features, Apple sets new consumer expectations for performance, comfort, and usability in MR devices.

Competitors and collaborators in the MR space will need to innovate to match or exceed the Vision Pro's capabilities, fostering technological advancements across the industry.

The Vision Pro's advancements in MR signal a shift toward devices that integrate seamlessly into users' lives, bridging the gap between digital and physical worlds in unprecedented ways.

Integration with Emerging Technologies

The Vision Pro's potential is amplified by its ability to integrate with emerging technologies such as AI, ML, and the **Internet of Things (IoT)**. These integrations enhance the device's capabilities, creating opportunities for more personalized, adaptive, and interconnected experiences.

Artificial Intelligence and Machine Learning

AI and ML play a critical role in optimizing the Vision Pro's performance and expanding its functionality. Apple is planning to add Apple Intelligence to visionOS 3 in 2025.

AI algorithms analyze user behavior and preferences to customize experiences. For example, the Vision Pro could suggest optimized layouts for virtual workspaces or recommend content based on past interactions.

ML enables the Vision Pro to adapt to its environment and user needs in real time, such as recognizing objects in the physical world and overlaying relevant digital information.

AI-powered voice recognition and NLP allow for seamless voice interactions, making it easier for users to navigate and control the device without manual input.

Integration of IoT

The Vision Pro's integration with IoT devices enhances its functionality by enabling connectivity across smart home and workplace environments.

The Vision Pro can interface with IoT devices such as smart lights, thermostats, and security systems, allowing users to control their environment through spatial interactions.

IoT sensors combined with the Vision Pro's capabilities can create context-aware environments. For example, the device could adjust lighting and temperature based on the user's activities or provide alerts when specific conditions are detected.

In professional settings, the Vision Pro can connect to IoT-enabled devices like interactive whiteboards and telepresence robots, facilitating more effective collaboration.

Expanding the Vision Pro's Capabilities

The integration of AI, ML, and IoT will enable developers to create increasingly sophisticated applications for the Vision Pro.

AI algorithms could monitor its hardware performance, predicting issues and optimizing battery life and processing power.

By combining IoT data with AI, the Vision Pro could create immersive educational experiences tailored to individual learners.

Its integration with IoT medical devices could revolutionize patient care by providing real-time diagnostics and virtual consultations.

By leveraging emerging technologies, it becomes a gateway to an interconnected digital ecosystem.

Impact on the Spatial Computing Ecosystem

The Vision Pro's introduction has far-reaching implications for the growth of the Spatial Computing ecosystem. Its capabilities and ecosystem integration encourage the development of interconnected 3D environments and lay the groundwork for the spatial web.

Growth of Interconnected 3D Environments

The Vision Pro fosters the creation of immersive 3D environments where users can interact with digital content in meaningful ways.

It supports shared virtual environments, enabling multiple users to collaborate on projects, attend meetings, or socialize in fully immersive settings.

The Vision Pro's success could drive the development of standards for Spatial Computing, ensuring compatibility across devices and platforms.

Vision Pro as a Foundational Step toward the “Spatial Web”

The Vision Pro represents an early but significant step toward the realization of the spatial web, a concept in which the internet evolves into a fully immersive 3D space.

Its ability to integrate AR and VR experiences could lead to the creation of persistent virtual environments where users can interact with content and each other in real time.

Just as the traditional web is built on hyperlinks, the spatial web will rely on interconnected 3D objects. The Vision Pro's support for spatial interactions positions it as a key enabler of this new paradigm.

By overlaying digital content onto physical spaces, the Vision Pro helps blur the boundaries between the real and virtual worlds, creating a seamless spatial web experience.

Accelerating Ecosystem Growth

The Vision Pro's integration into the Spatial Computing ecosystem will drive growth and innovation across industries.

Apple's tools and resources empower developers to contribute to the ecosystem, creating applications and experiences that attract more users.

As the Vision Pro gains traction, it will encourage consumers to explore and adopt Spatial Computing technologies, further expanding the ecosystem.

Partnerships between Apple, developers, and enterprises will foster innovation and create new use cases for Spatial Computing.

The Vision Pro's impact on the Spatial Computing ecosystem extends beyond its immediate capabilities, positioning it as a cornerstone of the next generation of digital experiences.

The Apple Vision Pro is not just a device—it is a glimpse into the future of Spatial Computing. By advancing MR, integrating with emerging technologies, and driving the growth of interconnected 3D environments, it lays the foundation for a new era of immersive experiences. Its potential to serve as a stepping stone toward the spatial web underscores its transformative impact on how people interact with digital content. As developers, businesses, and consumers embrace the Vision Pro, its influence will extend far beyond its initial release, shaping the evolution of technology and society in profound ways.

Case Studies and Early Use Cases of the Apple Vision Pro

The Apple Vision Pro has already begun to demonstrate its transformative potential through early use cases, developer insights, and pilot projects across various industries. While still in its infancy, its ecosystem has sparked innovation among developers, captured the interest of consumers, and initiated trials in industries such as healthcare, education, and entertainment.

Developer Insights: Building the Future of Spatial Applications

The Vision Pro's launch has catalyzed a wave of creativity and experimentation among developers, who are leveraging Apple's visionOS SDK and tools like Reality Composer Pro to create immersive applications.

Early Application Development

Developers across diverse fields are already building apps tailored to the Vision Pro's unique capabilities.

Companies like Zoom and Microsoft are developing spatial versions of their collaboration tools. For instance, Microsoft is working on a Vision Pro-optimized version of Teams, allowing participants to meet in immersive 3D environments and interact with shared virtual whiteboards.

Game developers, such as the VFX studio ILM, are creating MR experiences that leverage its advanced graphics and gesture controls. An early example includes immersive puzzle games where users manipulate 3D objects in their physical environment.

Medical software companies are developing surgical simulation apps that use the Vision Pro to provide real-time guidance and virtual practice environments for medical professionals.

Challenges and Opportunities for Developers

While the Vision Pro presents immense potential, developers also face challenges in creating applications for such a sophisticated platform.

Developing spatial applications requires mastering new tools and concepts, such as 3D modeling and spatial navigation, which can be a steep learning curve for traditional developers.

Ensuring smooth performance on a device as advanced as the Vision Pro demands high levels of optimization, particularly for graphics-intensive applications.

Despite the challenges, developers recognize the long-term opportunities of being early adopters, including access to a high-value user base and the chance to shape the future of Spatial Computing.

The Vision Pro's developer ecosystem is rapidly expanding, with early insights highlighting both the promise and the complexities of building for Spatial Computing.

Consumer Reactions: Early Feedback and User Experiences

As one of the most anticipated products in recent years, the Vision Pro has generated significant interest and curiosity among consumers. Early feedback provides a glimpse into how users perceive the device and its potential impact on daily life.

Positive Reception

Many early adopters have praised the Vision Pro's immersive capabilities and seamless integration with the Apple ecosystem.

Users have described its 3D visuals and spatial audio as unparalleled, creating a sense of presence that makes the content feel tangible.

The device's intuitive controls, including eye-tracking and gesture recognition, have received positive feedback for being natural and easy to learn.

Professionals who have tested the Vision Pro in work settings report increased productivity, citing the ability to organize multiple virtual screens and collaborate in immersive environments.

Concerns and Criticisms

Despite the positive reactions, consumers have also raised concerns about certain aspects of the Vision Pro.

Many consumers find the \$3,499 price tag prohibitive, limiting the device's appeal to a niche audience of early adopters and professionals.

Some users have reported discomfort during extended use, particularly due to the device's weight and reliance on an external battery pack.

The two-hour battery life is seen as a limitation for users who wish to use the device for longer sessions.

Anticipation for Future Updates

Consumers are optimistic about the Vision Pro's potential to evolve and address current limitations.

Many believe that future iterations of the Vision Pro, with reduced costs and improved ergonomics, will make it accessible to a wider audience.

Early users are excited to see how developers will continue to innovate, creating new applications that expand its functionality.

Consumer feedback highlights both the strengths and areas for improvement, providing valuable insights for Apple and developers as they refine the Vision Pro experience.

Industry Adoption: Pilot Projects and Early Use Cases

Beyond individual consumers, industries are exploring the Vision Pro's potential through pilot projects and early adoption efforts. These initiatives demonstrate how the device can transform workflows and create value in professional settings.

Healthcare: Transforming Training and Patient Care

The healthcare industry has been one of the earliest adopters of the Vision Pro, leveraging its capabilities for training, diagnostics, and treatment.

Hospitals are using it to simulate complex surgical procedures, allowing medical students and professionals to practice in a risk-free virtual environment. For example, a pilot project at a leading teaching hospital enabled surgeons to rehearse operations with highly detailed 3D models of patient anatomy.

Physical therapists are exploring its immersive capabilities to create engaging rehabilitation programs for patients recovering from injuries. Virtual environments can guide patients through exercises while providing real-time feedback.

The Vision Pro is being tested for telemedicine applications, allowing doctors to interact with patients in virtual spaces where medical data and diagnostic tools are seamlessly integrated.

Education: Immersive Learning Experiences

Educational institutions and EdTech companies are piloting the Vision Pro to enhance teaching and learning experiences.

Teachers are using it to create immersive lessons that bring subjects to life. For instance, history classes can take virtual field trips to ancient civilizations, while science students can explore the human body in 3D.

The Vision Pro is being tested as a tool to inspire interest in STEM fields. Virtual labs and interactive simulations make complex concepts more accessible and engaging for students.

Entertainment: Redefining Media Consumption

The Vision Pro is transforming how entertainment companies create and deliver content.

Music and event organizers are using it to stream live performances in 3D, offering viewers an immersive experience that feels like being at the venue.

Filmmakers are experimenting with interactive storytelling, allowing audiences to navigate 3D scenes and influence the narrative.

Corporate Environments: Productivity and Collaboration

Businesses are testing the Vision Pro for applications in productivity and remote collaboration.

Companies are piloting it to enable employees to set up virtual workstations with multiple screens and tools, regardless of their physical location.

Its ability to create shared virtual environments is being used to enhance remote meetings, enabling teams to brainstorm and collaborate on projects more effectively.

The Apple Vision Pro's early use cases and pilot projects underscore its transformative potential across a wide range of applications. Developers are creating innovative apps that leverage the device's unique capabilities, while consumers are providing valuable feedback that highlights its strengths and areas for improvement. Industries such as healthcare, education, and entertainment are leading the way in adopting the Vision Pro, showcasing how it can redefine workflows and create new opportunities.

As more developers and businesses embrace the Vision Pro, its ecosystem will continue to grow, unlocking even greater possibilities for Spatial Computing. These early insights and case studies lay the foundation for a future where it becomes an integral part of both professional and personal life, shaping the evolution of technology and society.

The Future of the Vision Pro

The Apple Vision Pro signifies a pivotal moment in the evolution of Spatial Computing, setting the stage for transformative changes in how people interact with technology and the digital world. By integrating advanced hardware, seamless software, and intuitive user interfaces, the Vision Pro redefines what MR devices can achieve. As we reflect on its impact, speculate on its future, and consider its call to action for developers and innovators, it becomes evident that the Vision Pro is not merely a device—it is the cornerstone of a new era of immersive technology.

Summarizing Vision Pro's Impact

The Vision Pro has already made waves across industries and among early adopters, highlighting its transformative potential for Spatial Computing. Its blend of AR and VR capabilities offers a new dimension of interaction that bridges the digital and physical worlds.

With features like micro-OLED displays, eye-tracking, and spatial audio, the Vision Pro sets a new benchmark for immersive experiences. Its ability to create fluid transitions between AR and VR environments has elevated the concept of MR to a level that feels intuitive and natural.

Its applications in entertainment, education, healthcare, and productivity demonstrate its versatility. From immersive gaming and virtual classrooms to surgical simulations and remote workspaces, the device opens up possibilities that were once confined to science fiction.

By leveraging visionOS and Apple's existing ecosystem, the Vision Pro has fostered a burgeoning community of developers and creators. This ecosystem is instrumental in driving the adoption of Spatial Computing and ensuring the Vision Pro's functionality continues to grow over time.

Its impact extends beyond its immediate features and capabilities. It represents a paradigm shift in how users engage with technology, creating opportunities for innovation and redefining expectations for immersive devices.

Looking Ahead: The Next Decade of Immersive Technology

The introduction of the Vision Pro marks the beginning of a new chapter for Spatial Computing, one that will shape the trajectory of immersive technology over the next decade. Its influence will be felt across industries, technological advancements, and societal adoption.

Paving the Way for Mainstream Adoption

While the Vision Pro currently targets early adopters and high-end users, its success will likely lead to the development of more affordable and accessible versions in the future.

As production scales and costs decrease, Spatial Computing devices will become more accessible to a wider audience. This democratization will drive mainstream adoption and integrate immersive technology into everyday life.

Advances in AI, ML, and IoT will further enhance the capabilities of Spatial Computing devices. The Vision Pro will likely act as a hub for these technologies, enabling smarter, more adaptive experiences.

Revolutionizing Work and Communication

The Vision Pro has the potential to redefine how people work, communicate, and collaborate.

The next decade will likely see the widespread adoption of virtual workspaces, enabling teams to collaborate across physical and digital boundaries with unprecedented efficiency.

Spatial Computing can bridge geographical divides, providing access to education, healthcare, and professional opportunities for underserved populations.

As the Vision Pro evolves and inspires future innovations, its introduction will be remembered as a turning point that reshaped the digital world and its interactions with the physical one.

Call to Action for Developers and Innovators

The Vision Pro's success—and the broader adoption of Spatial Computing—depends on the creativity and ingenuity of developers, innovators, and businesses. Apple has provided the foundation, but it is up to the broader tech community to explore its full potential.

Embrace the VisionOS Ecosystem

Developers have an unprecedented opportunity to pioneer applications that leverage the Vision Pro's capabilities. By embracing tools like Reality Composer Pro and visionOS SDKs, developers can create applications that redefine user experiences.

Innovators should prioritize creating applications that are inclusive and accessible, ensuring that Spatial Computing benefits people from all walks of life.

Explore Diverse Use Cases

The potential applications of the Vision Pro extend far beyond its initial use cases. Developers and businesses should explore new ways to integrate the device into industries and daily life.

The Vision Pro's immersive capabilities offer endless possibilities for storytelling, gaming, and interactive content. Creators should push the boundaries of what's possible in these fields.

Advocate for Ethical Development

As Spatial Computing technology becomes more integrated into society, ethical considerations must remain at the forefront of development.

Developers should prioritize protecting user data and ensuring transparency in how information is collected and used.

Innovators should consider the environmental impact of Spatial Computing devices and work toward sustainable design and production practices.

The Vision Pro represents a once-in-a-generation opportunity to shape the future of technology. By contributing to its ecosystem and championing responsible innovation, developers and innovators can ensure that Spatial Computing becomes a force for positive change.

The Apple Vision Pro's transformative impact on technology, creativity, and human interaction underscores its significance as a foundational device for the next era of immersive experiences. As we look ahead to the possibilities of the spatial web, interconnected environments, and groundbreaking applications, the Vision Pro stands as a beacon of innovation and potential.

To fully realize this potential, the collective effort of developers, businesses, and users will be essential. By embracing the Vision Pro and exploring its capabilities, we can build a future where technology seamlessly enhances our lives, blurring the boundaries between the physical and digital worlds. The journey has just begun, and the opportunities are limitless.

In our next and final chapter, we revisit how Spatial Computing affects us in more detail.

14

How Human?

Reaching business nirvana could trigger issues that verge on the unethical. When systems are set up to be hyper-vigilant in order to maximize information gathering and understanding, it is not surprising that boundaries that should not be crossed are. Given that Spatial Computing is in its relative infancy, those boundaries need to be identified. Here, we first provide a framework that will aid in understanding why human beings tend to have cognitive issues with radical new technologies, and then we provide a background on current issues relating to privacy, security, identity, and ownership, and identify their relevance to Spatial Computing. We then proceed to discuss how Spatial Computing business nirvana and human social good could intersect as much as possible.

It does not escape us that in the process of creating and using technologies to aid and entertain, human beings can become slaves to those very technologies, which have the effect of dehumanizing them. Another passage to dehumanization is the effort to reduce too easily to data the characteristics of human individuals and human life for the sole purpose of commerce, rendering humans down to basic tools of utility. When human beings overly view themselves and others as consumption machines and the content of their lives as data points, they become estranged from themselves. An inherent contradiction is that this kind of rationalization is in itself very human. The key to avoid estrangement is to be balanced between our actual human needs and human wants. We have it in our power to manage that balance when it comes to Spatial Computing. Before we can do so, we need to have a good understanding of Spatial Computing technologies, which can only come out of experiencing them.

Transformative Experience

In her book *Transformative Experience*, Yale philosopher L.A. Paul makes the point that if someone has not experienced something that is hard to imagine due to lack of similar experiences to compare to, then it is very difficult to predict what that experience would be like. The only remedy to this would be to go through the experience.

Opinions about what the experience would be like, versus what the experience is actually like for the person who finally undergoes it, usually do not converge. The connections to Spatial Computing technologies, and how it is difficult to understand them without experiencing them, is the point here. People who have not experienced Virtual Reality and/or Augmented Reality really do not have a good understanding of it, though opinions abound.

Similarly, fully understanding the benefits and utility of Spatial Computing is not possible without actually using Spatial Computing technologies. Many people have issues having to do with digital privacy and are anticipating that there will be many more problems surrounding privacy in the future. They are not wrong about this. Spatial Computing brings with it a whole slew of possible problems here that could be akin to opening up Pandora's box. However, as we humans become more understanding of what Spatial Computing is good for, we feel that there will be a mediation and balance between individual utility and a person's boundary to digital privacy.

This boundary was severely tested recently, most notably in the Facebook-Cambridge Analytica scandal. We review the scandal in the following section, along with some other examples, in order to show what the current environment is like with regard to digital privacy.

Privacy and Security

Let's now think further about Spatial Computing's impact upon privacy and security. By understanding how the events that follow unfolded, and where the issue regarding privacy really lies, it will be easier to work through how issues regarding privacy and Spatial Computing should be addressed.

On Facebook and Cambridge Analytica

You might have heard about Cambridge Analytica, as many stories broke about it in 2018. There was a huge outcry raised by revelations that it had gathered data from millions of Facebook users, often without those users' knowledge or permission.

It then used that data to build a new kind of Facebook advertising, aimed at manipulating citizens to change their voting behavior. The details tell a cautionary tale as we head into a Spatial Computing world where even more powerful data collection and manipulation of users is possible.

The Facebook-Cambridge Analytica scandal started in June 2014, when a researcher named Aleksandr Kogan developed a personality quiz app for doing psychometric research with Facebook users, called “This Is Your Digital Life,” as part of the research that he was doing at the Psychometrics Centre at Cambridge University. This personality quiz app appeared to be very innocent, with content reshared by millions of people around the world. The Psychometric Centre had been doing a lot of social media research and had created several apps. Those apps had already shown that Facebook users could be manipulated with various content techniques that the lab had developed, and that the data collected from Facebook could show all sorts of sensitive information about people’s personalities and political persuasions.

The earlier apps were built by John Rust and his team. The goal of that research, and indeed his entire 40 years spent in the field, was to warn the world about what could be done with that data and the dangers of allowing it to be so freely traded. Rust had warned the University that Kogan’s app was taking the research done in the lab further than Rust was comfortable with. Kogan’s app slurped up the data from millions of Facebook users, most of whom didn’t give approvals to the app for their data to be collected.

Most users didn’t know that when you give these apps permission to gather data that the apps get access to your entire network of friends, which Facebook calls “the social graph.” Facebook is built around this social graph, which is like a rolodex, but with a lot more things—from photos, posts, and other data—than a simple rolodex would have in it. The app collected the same data on each of your friends, which meant that the database could quickly grow, and it did—to 87 million users by the time it got sold to a new start-up named Cambridge Analytica. Most of those app users had no idea that much of their data now sat in Kogan’s databases.

As the developer on this app, Kogan could access data about the users who downloaded it and used it, along with the data of their friends too. “It would go into their entire friend network and pull out all of their friends’ data as well,” says Christopher Wylie, one of the founders of Cambridge Analytica—now a whistleblower. “The company itself was founded by using Facebook data.”

The collection of data on so many people without their knowledge is one thing, but what Cambridge Analytica did with that data was a whole other matter. Kogan learned a whole lot about how to use psychometrics to manipulate voters.

Psychometrics is the science of measuring people's mental and psychological traits, strengths, and weaknesses. When he sold the data to Cambridge Analytica, he also helped the company use the user manipulation techniques that the lab had also been studying.

When Kogan's app asked for that data, it saved that information into a private database instead of immediately deleting it. Kogan provided that private database, containing information on about 50 million Facebook users, to the voter-profiling company Cambridge Analytica, which used it to make 30 million psychographic profiles about voters. This was laid out in several articles by the New York Times, the Wall Street Journal, the Guardian, and Robinson Meyer in the Atlantic.

“We would know what kinds of messaging you will be susceptible to and where you’re going to consume that. And then how many times you need to touch that to change how you think about something. It is a full-service propaganda machine,” Wylie has said.

The Cambridge Analytica scandal showed Facebook (now Meta) had been playing fast and loose with private data, and this saw its stock price drop by \$100 billion in value in the space of just one day—at that time about 12 percent of its value—on March 26, 2018 (although it subsequently bounced back). It wasn’t the only thing that caused market outrage at Facebook, either, but it was the one that finally got the public to understand that Facebook shared data with more than just itself or other users on the system, and that knowledge dramatically reduced trust in the system and encouraged a new movement.

For a long time after the Cambridge Analytica scandal, the cry “delete Facebook” was alive on blogs, in media, and on Twitter. People including Cher, Elon Musk, Will Ferrell, WhatsApp cofounder Brian Acton, Rosie O’Donnell, Apple cofounder Steve Wozniak, and tech journalist Walt Mossberg, among others, announced they had deleted their Facebook accounts and urged others to do the same in tweets with a #DeleteFacebook hashtag. #DeleteFacebook people often reminded you on Tweets, and elsewhere, that Facebook follows you around the internet, building a dossier on you so that it can sell it to advertisers using targeting techniques, and that the service watches everything you do to help make its advertising system more efficient. In another privacy scandal in 2018, Facebook shared data with other companies like Amazon, Spotify, Netflix, and Bing, giving these companies special access to users’ data without anyone else knowing. Why? To get even more data on users so that Facebook’s ad targeting system would work even better.

But why does this all cause outrage? In our research, many people really hate being digitally targeted by ads. Many tell us about doing searches for something, like a ski trip, and then ads for skis or four-wheel-drive cars follow you around no matter where you go on the web.

Lots of people don't like this kind of targeted advertising, nor do they like that it follows you in a process that the industry calls "retargeting."

It's one thing to have a few corporations following you around as you move around the web, pushing ads in your face. To some that actually might be helpful, if it were something you needed, although invariably the process isn't perfect so you get ads for a car or coffeemaker you just bought, which is both freaky that the system knew you were just showing interest in something like that, and sad because it didn't get it right. However, when similar techniques are used for political manipulation, it brings a whole new level of outrage because the act degrades our democracy and our own part in choosing who we want to represent us.

Then there are a variety of other concerns about Facebook as well, including that it has been built to addict users and take advantage of how their brains work. Salesforce CEO Marc Benioff even said "You can see Facebook is the new cigarettes for our society. It's something that badly needs to be regulated. They're certainly not exactly about truth in advertising. Even they have said that. That's why we're really in a crisis of trust." Academics back up that claim: a study published in November 2019 from Stanford University and New York University even found that quitting Facebook is good for your mental health (<http://web.stanford.edu/~gentzkow/research/facebook.pdf>).

With Spatial Computing in the offing, especially when it comes to devices we will wear on our eyes that will bring us some mixture of analog and Virtual Reality, these issues become even more amplified. Soon we might see virtual ads on every surface in our lives, and we will have visualizations that can bend what we think of as real, and that present powerful new challenges for the human brain, from potential new addiction problems to inabilities to tell the virtual from the real.

We have done a lot of thinking about privacy when it comes to Spatial Computing, and here are a few of the issues that will come up as we move to living more and more of our lives in some form of digital reality, along with new ways of living, assisted by autonomous vehicles, robots, and synthetic beings.

In the future, it will be possible for big companies and governments to collect a much more complete set of data about you. Where today's systems have to read your comments, follow you around with location-tracking systems, or watch what you like or follow, in our Spatial Computing future these same companies and governments will know a lot about your health. Collecting your heart rate and DNA is just the start. Imagine a set of companies knowing everything you touch, look at, or eat.

And, in the Spatial Computing world, these companies and governments could know a lot more about every social interaction you ever have. Meet someone in the street? It takes note. Go on a date with someone? Play a game? Make a video call? All of this could potentially get tracked in great detail. Eventually the system could deeply know who your business associates, friends, lovers, and infrequently met friends are, your exact activities, and the content of your life.

Furthermore, in the future you will be wearing devices that can present such compelling visuals and audio that present new powerful capabilities. In our other chapters, we discussed how VR is being used to treat a variety of brain ailments, from PTSD to dementia. With such powerful positive capabilities also comes some negative ones. These could be used to manipulate people in a much more complete way than how Cambridge Analytica did with Facebook ads.

There are some choices that we'll have to make over the next decade, as we start to live more of our lives with digital interfaces all around: do we want to have every action we take studied, particularly when we go into stores? Already Amazon Go stores watch literally every step, every touch, every look we take at products. That lets Amazon build powerful new databases about our lifestyles and brand choices that can reveal things about ourselves we didn't even know. Target, the New York Times reported, knew of one young woman's pregnancy before she told her father, based on her buying behavior. Are we ready for other such revelations?

Another choice we'll have to make is associated with advertising. Advertising, at its core, is trying to urge you to take actions you didn't know you wanted to make. "Drink Coke," an ad urges, and does so in compelling ways and at compelling times, like when you are sitting down in a movie theater.

Will we choose systems that have such deep advertising links into every part of our lives? Will there be any control or choice? For instance, instead of seeing a virtual Coke ad at the beginning of a movie, will we have the control to say "please don't show me any more Coke ads; replace them with messages reminding me to drink more water instead, or, don't show me any ad at all"?

Finally, we soon will have virtual transactions on top of both real things ("hey robot, can you bring me a hamburger") and virtual ones, like buying a virtual costume to wear in a virtual world, or a virtual boat ride—something that's already possible in one of the Virtual Reality metaverses we've been in. Whenever we carry out one of these transactions, do we want corporations and governments to know about it? Well, governments will be hard to keep out of these things, but Apple will certainly try to become the main payment processor with Apple Pay. It won't be alone trying to become the main way we buy things.

Of course, we will use systems that do all this stuff, particularly if they provide great utility in return, but as we start to use Spatial Computing glasses and use them to order all sorts of things, from hamburgers, to music, to new kinds of entertainment, and as we use them to control and visualize a new world with robots and autonomous vehicles, the consequences of giving up our data will be discussed, challenged, and questioned even more.

Let's check out another example that shows the power of the data we are constantly generating.

OJ Simpson and Location Data

OJ Simpson was a fugitive from justice on June 17, 1994. The police had an arrest warrant for him, as he was a suspect in the murder of his ex-wife Nicole Brown Simpson and her friend Ron Goldman. On that day, Al Cowlings, OJ Simpson's close friend, was driving his white Ford Bronco on a number of Los Angeles freeways, with OJ Simpson in the passenger seat. Cowlings reportedly made a call to 911, saying that Simpson was armed, with a gun pointed to his own head, demanding that Cowlings drive him to Simpson's estate in Brentwood or Simpson would kill himself. That call was traced to the Santa Ana Freeway. Cowlings made several other calls that were also traced, and police officer Ruth Dixon caught up with them northbound on Interstate 405.

How was his car found? Cell phone tower triangulation. The way cell phones work is they broadcast their location to towers nearby. One to three towers picks up the signal and the system itself tracks each phone and hands it off to new towers or "cells" as its user drives around. This data can't be encrypted, because if it were, the system the way it is currently set up wouldn't be able to effectively work. The triangulation that caught OJ Simpson was accurate enough to focus police attention on a single freeway.

These days, police have even more location data to use, since each smartphone not only uses cell phone triangulation, which isn't very accurate, but has a GPS chip inside that reports its location via services and apps. Apple, for instance, tracks cell phone location for its "Find My iPhone" service, and many mapping apps record your location, as well. Some, like Verizon, even sell real-time data feeds that are discerned from this data to other companies like Google, which uses this data to figure out traffic conditions. While that data is cleaned of personally identifiable information, it could be discerned who is associated with that data later on by piecing together other data about users.

Soon, with Spatial Computing glasses on many people, and autonomous cars and robots rolling around, police will have even more tools to find people.

Apple and the FBI

It started with a 2015 terrorist shooting in San Bernardino and ended with the FBI filing a lawsuit to get Apple to unlock an iPhone owned by a suspect. Fourteen people were killed during that shooting, and the FBI wanted more information about the shooters and who else they could be connected with.

After the shooting was over, several phones were collected, most of which had been destroyed by the shooters. One that had survived, an iPhone 5C, was studied deeply in an effort to find information about what was behind the killings. Apple was then served with a court order to turn over information from the phone and account of the owner, which it did. Apple worked with the FBI before the lawsuit was filed and helped it uncover a bunch of metadata from encrypted files on the phone, along with what they could get from the cloud account that was associated with the phone.

The FBI, though, soon wanted more. It wanted Apple to unlock the phone itself, which would let the FBI look at all the data and files on the phone. Apple actually did work with the FBI, who tried to get Apple to back up the data and files because backups aren't encrypted, but the officials who were studying the phone had already taken steps to reset the iCloud password. It turns out that if you reset an iCloud password, the only way for a phone to back up is to put a new password in the phone. However, a new password will not be accepted if passcode protection was enabled by the phone's owner, and this is what happened in this case.

Months later, the FBI filed a lawsuit to attempt to get Apple to unlock the phone or add a backdoor that law enforcement could use to force a phone to unlock and decrypt any encrypted data stored inside. Apple's executives had already met to decide what to do if Apple was ever asked to provide a backdoor. Obviously, Apple wanted to provide any data it could to help catch and stop more terrorist acts and knew that the public might not understand a nuanced argument about backdoors, encryption, and privacy rights, but Apple had decided to hold the line when it came to the phone unlock itself.

Tim Cook, Apple's CEO, ardently argued that civil liberties were at threat with the ask from the FBI, and said about it, "So the act of banning or limiting or putting a back door in—a back door is any vulnerability that if I can get in, that means somebody else can... No one should have a key that turns a billion locks. It shouldn't exist." Cook felt that giving up privacy for security meant exposing "99 percent of good people." The FBI withdrew its case after government officials reported a third party had managed to unlock the iPhone, leaving the case without a clear legal precedent.

This issue will keep coming up, especially when we move more of our computing to Spatial Computing glasses. More of our data will be encrypted and, sometimes, police and other government officials will desperately want to unencrypt that data to catch terrorists, or other criminals. Apple is making the strongest defense of privacy as we move into this new age and it'll be interesting to see if it can keep holding the line.

Apple was facing different demands from the public and the government due to COVID-19 that show its privacy philosophies in a different light. In April 2020, it announced that, along with Google, Apple would be adding a tracking system into its iOS operating system that runs inside every modern iPhone. This system would track COVID-19 spread, allowing users to share data through Bluetooth Low Energy transmissions and via approved apps from health organizations. This allowed phones to share lots of data with other phones that were in close proximity to each other. Official apps from public health authorities got access to this data, and users who downloaded them could then report if they've been diagnosed with COVID-19.

Privacy has long been balanced by security. A parent whose child has been kidnapped would gladly hand over all sorts of personal information in an attempt to get the child back. Sometimes, the privacy ask is too big and doesn't represent enough of a public good to be worth the trade-off, and hopefully these examples have demonstrated some of those decision points. Next, we'll look at identity and ownership, which will back up the trade-offs made in privacy and security.

Identity and Ownership

To have dependable privacy and security, identity and ownership roles need to be defined strongly in tech systems and apps. It's like a building. If you want strong security or high privacy, the identity of everyone who enters that building needs to be known, as well as the ownership rights, or usage rights, for access points in that building. Here we discuss some areas where identity and ownership discussions have changed recently due to the growing use of particular radical new technology.

Blurred Faces on Google Street View

You might have seen a Google Maps car roll by. It has a ball of cameras on top of a car that's being driven around, capturing images and taking other data for its maps and Google Street View. While it's legal to capture images of people in public streets, these new captures have raised a number of new privacy debates. After all, most of our public photography laws were written during a time when capturing images of people required an act of pushing a button on a camera, not driving down the street slurping up everything the camera sees.

Back in 2009, Privacy International's Simon Davies argued that Google should have obtained consent from everyone who is featured in such photos, and so have various other people and governments, including Switzerland's data protection commissioner, Hanspeter Thüer. Most of those arguments have gotten thrown out by courts or fixed by technology that is much better at blurring out people's faces, along with other identifiable information, like car license plates, which Google has done for its Street View.

In the United States, efforts to get these cameras banned has also been refused because US law allows photography in places where there isn't a clear expectation of privacy. In a world of robots, autonomous cars, and people walking around with Spatial Computing glasses, we are going to see a public street that is mapped out in real time and in great detail.

Already, this is happening because of self-driving vehicles. Zoox, Waymo, Tesla, and others have already shown off new "high-definition maps" that have extremely high precision, with everything imaged down to centimeter level, from signs to potholes, to barriers. As newer cars drive by you, they are also watching you, categorizing you as a pedestrian, bicyclist, runner, and even getting more details like gender and age, which can be used to predict your next likely action. A child is more likely to run quickly into a street than, say, an 80-year-old man would.

Such a car also needs to identify certain people differently than others. Police officers, garbage truck workers, road crews—all will be tracked by cars rolling by. There's a good chance that lots of others will be caught by the cameras on these cars and broadcast to others, too, and it's not clear that their faces will always be blurred. Already we are seeing Teslas capturing other cars on roads and broadcasting their images, along with people walking next to the car, or, worse, criminally damaging the car. This could lead to new privacy fights.

As Spatial Computing glasses get added to the mix, having similar technology to self-driving and mapping vehicles, we'll see new databases being made of people on the streets, leading to new concerns about privacy. At Meta, they are planning for a day when your glasses will recognize a friend you meet, and will show you things like their last post, or present the opportunity to play a new kind of game with them, too.

China and Facial Recognition

People who have never been to China have little idea just how much more aggressively China uses technology than other countries. Citizens no longer use credit cards. Everything is purchasable via WeChat; in America, many businesses don't yet accept Apple Pay. In the cities of China, even street vendors take payments via WeChat and no one carries cash anymore.

The Chinese government had aimed to have even rural villages go cashless by the end of 2020, which still has not come to fruition. It is with the uses of facial recognition and Artificial Intelligence, though, that the differences in technology usage between China and other countries is even more extreme. While in the United States, Facebook has had facial recognition technology since its \$100 million purchase of Face.com in 2012, uses outside of China have been fairly muted, and there's a resistance to facial recognition in many places that simply doesn't exist in China. Cities like San Francisco have even banned its use by governmental agencies.

Within the next 10 to 15 years, though, millions of cameras are forecast to be moving around the large cities in the world via autonomous vehicles. Some autonomous vehicles, like those from Waymo, have 19 cameras, not to mention other sensors, and during the 2030s, we should see tracking capabilities that will make what's going on in China today seem pretty quaint.

Many Westerners might key in on the human rights abuses that come with China's use of this technology, and these abuses are worth focusing attention on—whether it is its harsh crackdown on ethnic Muslims in its Western region, its forced hospitalization of people who were tested with an elevated body temperature as the novel coronavirus outbreak started being taken seriously, or the tracking of literally every move of citizens and the building of a social score of each. But it is the range of uses that we are focused on here. Everything from access controls at schools, factories, airports, and other places, to even using facial recognition in soda vending machines to enable quicker transactions.

Hundreds of millions of cameras are used in this facial recognition system throughout China, and the Chinese government forces its citizens to register their face when they pick up their phones, making for easy facial capture and registration. During the novel coronavirus outbreak, new uses came to the public's attention where citizens are tracked via apps on their phones, facial recognition, and other Spatial Computing technology.

One Chinese Deep Learning software company that makes many of the facial recognition systems China uses, Megvii, has filed to go public and has contracts to run cameras in public housing, where its service is used to curb subletting. The company has 2,300 employees and is valued at more than \$4 billion. Its product does face detection, face comparing, and face searching. It also goes beyond facial recognition, with human body detection, face merging, and photo clustering, where after recognizing faces of photos in an album, photos can be clustered based on faces automatically. It beat Google, Meta, and Microsoft in tests of image recognition at the International Conference on Computer Vision, with its product Face++, which is the world's largest open source Computer Vision platform. Megvii has an R&D lab in Redmond, Washington, near Microsoft's campus.

With its competitors SenseTime, CloudWalk, and Yitu, Megvii has brought facial recognition even into Chinese classrooms. For instance, China Pharmaceutical University uses facial recognition for everything from students' attendance to monitoring of students' in-class performance. Its use in classrooms is controversial, even in China, but some people think the system can help prevent dangers, reduce campus bullying, and improve teaching quality. One such system scans classrooms at Hangzhou No. 11 High School and records students' facial expressions, categorizing them as happy, angry, fearful, confused, or upset. It also records student actions such as writing, reading, raising a hand, and sleeping at a desk, even tracking them outside of classrooms—enabling the replacement of ID cards and wallets at the library and canteen.

Facial recognition in China is even being used to catch jaywalkers, but it does have flaws while doing this. In 2018, one such system caught famous Chinese businesswoman Dong Mingzhu (she made first place on Forbes list of the top 100 outstanding businesswomen in China) and put her face on a public screen dedicated to “naming and shaming” jaywalkers caught by the city of Ningbo’s facial recognition system. The problem was that she never was there—the system had simply caught an image of her that had appeared on the side of a bus. Similar systems in the city of Shenzhen are also set up to send an automatic notice of a fine, along with visually shaming jaywalkers on large LED screens at intersections, to the caught jaywalker’s WeChat, of course, where they can pay the fine right from their phone. How many infractions are there? Since 2018, there have been tens of thousands of jaywalkers in major cities like Beijing, Shanghai, and Shenzhen.

Issues with Facial Recognition in the US and Europe

Looking backward, we see that China, the United States, and Europe have very different approaches to privacy. This can be seen by traveling to the three. In China, you see a very tech-centric approach literally everywhere, as citizens use WeChat to pay for everything, and on many streets, grids of cameras watch everything. The use of smartphones is heavier in streets and on public transport than you’ll see elsewhere, too.

In China, citizens don’t have the right to freedom of speech like United States citizens do, and every word posted online is read by either a person or some Machine Learning algorithmic software that decides whether the posts can stay up or not. Behind the scenes, China has blocked most outside web services. Meta, Google, Twitter, and LinkedIn are usually blocked for citizens there, and it has developed its own versions of each, or in the case of WeChat, a service that really does the same as several of the services in one: shopping, payments, a social network, and a health tracker all in one app. This is backed up in China by extensive use of facial recognition and surveillance tracking.

Citizens know they are being tracked literally everywhere, and Machine Learning algorithms even give each citizen a “social score,” so if citizens post things that aren’t allowed, or perform things that the government doesn’t want done, even like jaywalking a street, their social score goes down. If it goes down low enough, they can be blocked from services, or even the ability to use public transit or take domestic or international flights.

Such a set of systems would be impossible to imagine in the United States, and particularly Europe, where the EU has written laws that protect the privacy rights of citizens. Travel through Germany, too, and you see a much higher percentage of people who still read newspapers, rather than the mobile phone usage you see in China, and you can see a different set of attitudes and adoption rates toward technology.

On the other hand, Europe is ahead of the United States in using technology to transact. Apple Pay, and other mobile-based payments, work literally everywhere in Europe, due to Europe’s much more ubiquitous adoption of contactless payment systems by even street vendors, taxis, and restaurants. The United States has not seen usage by most of these kinds of vendors yet.

These differences in transactional affordances, in the law, and in penalties, set us up for making some predictions about how much resistance each legal system will have to Spatial Computing devices, which include various forms of smart glasses, along with robots, drones, and vehicles that use Computer Vision, sensors, and AI algorithms for moving, or being moved, around the real world.

We can imagine that a person wearing glasses with several cameras (Qualcomm has shown a reference design with seven), controlling robots, and self-driving vehicles will be more immediately accepted in China than in Germany, and developers who make the systems that run on each will need to have different systems for different laws. Already, Tesla’s autopilot systems turn off features in the EU due to stricter laws than they have in, say, California. We can see one day, coming within a few years, where in Beijing and Shanghai a large percentage of the population will be wearing Spatial Computing devices almost everywhere, yet in Munich or London, we’ll see much lighter usage, and much stricter laws for how devices communicate in terms of, say, them recording or broadcasting.

We talked with Avi Bar-Zeev, Microsoft HoloLens co-inventor and former senior manager of AR prototyping for Apple, on what the reception might be to people wearing AR glasses in the near future. He believes that there should be an opt-in system for people that is relevant to whether or not someone has the right to take a photo or video of you while you are on a public street. While we don’t feel as strongly as him on this point, we do recognize that there are others that probably do.

One privacy law worth digging into, due to the fact that it will impact the way in which Spatial Computing sees adoption over the next decade, is the GDPR, or General Data Protection Regulation, which the EU started using in 2018. The GDPR says that your privacy is a right that will be protected by law, and that corporations and others that collect your data need to do so lawfully, with fairness and transparency. To be compliant with all of this, corporations need to get consent from users, notify them of breaches, let them access their data, let them delete it, (or have the system “forget” them), let them take their data elsewhere or be “portable,” and have privacy deeply built into their systems.

Luciano Floridi, Founding Director of the Digital Ethics Center at Yale University, who has consulted governments on digital privacy, told us that Europeans are adapting very well to GDPR, however, more work needs to be done to make future regulations flexible to allow for continuing European technical innovation.

The United States allowed a far more open system to collecting data, providing very few laws before GDPR when it came to the collecting of, and playing fast and loose with, privacy. This allowed companies like Meta, Google, Amazon, Netflix, and others to get going and to build massive databases of people’s data, from what they buy to what they say online in social networks, to what they watch in the evening on TV. That, in return, enabled the creation of massive new targeted advertising systems, particularly at Google and Facebook, that got built before the GDPR was put in place in 2018.

Today we are viewing a new set of challenges to privacy that even the GDPR doesn’t protect against, with cars that have neural networks and dozens of cameras and sensors that can slurp up all sorts of information, including whether you put your garbage cans out on the curb or not. Tesla added that capability to its entire fleet in 2019. We’re already seeing facial recognition being adopted, even in the more privacy-conservative Europe. It’s being used in shopping malls, supermarkets, airports, and by police who have millions of cameras to monitor. One French city, Nice, even adopted it citywide. This adoption needs to be compared to the legislation that came with GDPR, because GDPR set the tone for other laws in other countries. In the United States, California followed the GDPR with its own set of laws—the **California Consumer Privacy Act (CCPA)**, which is comparatively much more lenient in terms of company fines for privacy infractions.

GDPR and CCPA lay the beginnings of a new regulatory framework that will be needed when it comes to new Spatial Computing technologies. This is true whether we are discussing glasses or various other head-mounted displays with their array of cameras and sensors, or other Spatial Computing technologies, including autonomous vehicles and robots, all of which gather massive amounts of data on our streets, in new virtual metaverses, AR Clouds, and with virtual beings that will interact with us.

These earlier laws will still be needed, but with new ones tacked on to cover these new technologies. Necessary laws would cover what can be done with facial, and other identity recognition, and data from new sensors, whether ones that watch our eyes, or 3D sensors that can map out our homes, offices, and other places in great detail.

This gap became apparent as we researched dozens of cases where live facial recognition is used, both in the United States and in Europe, and we picked a few that demonstrate how the technology is being used across each continent. In the United States, we see use cases including police preventing human trafficking. Police departments tend to be big proponents of such technology, both in the United States and in Europe, because of the power of facial recognition to pick faces out of the crowd. Police camera manufacturer Wolfcom has contracts with 1,500 organizations and is pitching body cameras with live facial recognition to them. Competitor Axon, the largest manufacturer of body cameras in the United States, declared in 2019 that it would not put the invasive technology in its hardware, citing “serious ethical concerns.” Meanwhile California, Oregon, and New Hampshire have outlawed the use of facial recognition technology in body cams.

This use case won’t stay only in police hands for long. Meta has an internal strategy guide for its future Spatial Computing plans where it has drawings of how friends will recognize others in the street by wearing glasses with cameras on them, and some form of facial recognition plays a big role in its plans to enable new kinds of large-scale games that you can play with others in parks.

For now, however, we wait for these more fun uses of facial recognition and note that it was used to identify someone, later identified as Larry Griffin, who left a package that looked like a bomb in a New York subway in August 2019. His name and photos were sent to the cell phones of every cop in New York and he was arrested in the Bronx the same day and charged with three counts of planting a false bomb.

Facial recognition was also used in a variety of other police or security cases, from identifying the man charged in the deadly shooting at the Capital Gazette’s newsroom, to Taylor Swift using facial recognition to identify stalkers. A kiosk, it reported, was set up for fans to view videos of her. Inside the kiosk were cameras that took photos of those fans, then sent the images back to an office in Nashville where the images were cross-referenced with a database of people who had been identified as potential stalkers of the pop star.

Across the Atlantic Ocean, uses are similar. Police at Scotland Yard, the headquarters of London’s Metropolitan Police Service, began using live facial recognition technology in London to help identify criminal suspects in 2020.

The Metropolitan Police has claimed that 80 percent of people surveyed backed the move, the Guardian news organization reported, rejecting claims that the scheme was a “breathtaking assault on rights.” Nearby, in South Wales, police already use live facial recognition, and judges have ruled in favor of the technology.

Another example of the use of facial recognition in Europe is in retail stores. A Budgens supermarket in Aylesbury, Buckinghamshire, UK, had problems with shoplifting. The manager installed a facial recognition system called Facewatch, which warns him if a “subject of interest” walks into the store. If that happens, a manager asks them to leave. The manager there says the technology makes life easier. When he opened his shop, he would have things stolen every day or two, but since he introduced the system, it’s become less common. Plus, his staff feels safer, he said, especially late at night, which is good for team morale.

Airports around the world are seeing facial recognition systems pop up. One of Europe’s busiest airports, Amsterdam’s Schiphol, has tested facial recognition boarding. Soon passengers will be able to walk through customs and other security areas without pulling out their passports, thanks to the technology. We’ve seen similar systems show up in San Francisco, and in Dubai.

As we head into the Spatial Computing age, with AR headsets, autonomous vehicles, and robots capturing and receiving massive amounts of data, with an array of up to 30 cameras and sensors on each, we’ll see new capabilities for governments and companies to watch and track us than ever before.

Human Social Good

How do we, as responsible human beings, navigate between protecting our digital privacy and garnering the benefits that could be received by sharing data, both on an individual and wider societal level?

This question cannot be answered in a general way; the actual context in which data is being asked for is extremely relevant to answering it. And then, operationally speaking, it will take some kind of body of peers to ensure what has been decided would be appropriate for human social good.

We spoke with Dan Ariely, James B. Duke Professor of Behavioral Economics at Duke University, about these very issues. The way to think about human decision-making, from his perspective, is that human beings are not rational from a practical and behavioral perspective, but rather irrational, rife with biases and blind spots. Human beings tend to have irrational fears that are then used to make decisions. Because of this, beneficial opportunities are lost—opportunities that not only could benefit an individual, but rather a whole society.

An example that he used was that when a person visits a doctor with an ailment and the doctor asks that person questions about their health, there is very little worry about answering the doctor's questions, in effect, providing the doctor with data.

Exchange that with having a way to anonymize massive amounts of health data in order to find a cure for a disease and a significant amount of people have a problem with agreeing to provide their health data for this outcome. How is all of this connected to Spatial Computing?

With Spatial Computing, the amount of different kinds of data—from audio to textual to visual—that could be collected is enormous. In the wrong hands, that kind of information could be used in hacks to destroy whole communities and societies. Greedy, unrestrained companies could exploit that data for profit and destroy a person's piece of mind. However, in the right hands for the right purposes, benefits to individuals and societies could truly be vast and transformational.

When we asked the futurist Jason Silva what social project would be at the top of his list with which massive amounts of analyzed data could greatly be addressed, he very quickly answered with, "climate change." Brad Templeton, Futurist-in-chief at Electronic Frontier Foundation and Chair of Networks and Computers at Singularity University, told us that traffic patterns in large cities could be changed enormously for the better if there was some kind of massive organization of and analyses of large amounts of relevant data. But, governments tend to lag when it comes to investing in technology for societal change. The time lag between the wonderful things that are possible to be done at this time and when a very small portion of those things are put in motion to be done is very well known to anyone who has worked in politics and/or a very large bureaucratic organization.

So, how do we, as human beings, figure out what projects should be addressed with the use of data and then how to get them done? Ariely believes that it is best done by getting the top minds and experts in the relevant disciplines together to serve in societal peer review boards. They would produce recommendations to people as to how they should view company and organizational requests for their data and aid them in their decision-making in this regard. There are many issues that arise, ranging from the organizational to the practical, to how effective advice would be versus legal enforcement. All people, whether they are good with technology or not, will increasingly need to be able to make decisions about what they should do with their data. The true age of personalization is here, and with Spatial Computing it will become even more pronounced and perfected.

Along with Brian Green, Director of Technology Ethics at the Markkula Center for Applied Ethics at Santa Clara University, we foresee that beyond governments and other organizations getting involved in Spatial Computing privacy and ownership issues, religious leaders and religious orders will as well—with the Roman Catholic Church joining forces in February 2020 with Microsoft and IBM to collaborate on ethical “human-centered” ways of designing AI.

Somewhere in the farther future, when the dignity and rights of human-like robots becomes an issue, human beings can look back to see if the groundwork for Spatial Computing conventions, laws, and opportunities was laid down properly for human social good. It may be that at that time, robots do achieve consciousness as David Chalmers, Professor of Philosophy at New York University, believes they could. We had better start preparing for that possibility now as Spatial Computing grows in sophistication and human usefulness.

The choices we make today will define the world we live in tomorrow. Spatial Computing comes with both incredible potential and serious risks. It could reduce us to mere data points which are monetized, tracked, and categorized for someone else’s gain.

Technology itself isn’t the problem—it’s how we choose to use it. The challenge is to make sure that, in the process, we don’t lose sight of what makes us human.



packt.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

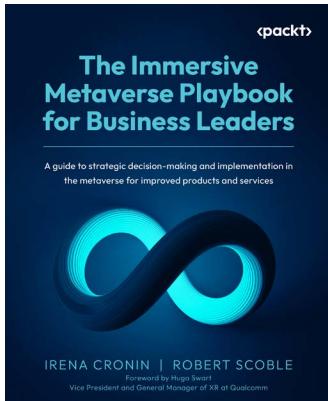
Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

At www.packt.com, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

Other Books You May Enjoy

If you enjoyed this book, you may be interested in these other books by Packt:



The Immersive Metaverse Playbook for Business Leaders

Irena Cronin, Robert Scoble

ISBN: 9781837632848

- Get to grips with the concept of the metaverse, its origin, and its present state
- Understand how AR and VR strategically fit into the metaverse
- Delve into core technologies that power the metaverse
- Dig into use cases that enable finer strategic decision-making
- Understand the benefits and possible dangers of the metaverse
- Plan further ahead by understanding the future of the metaverse

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit authors.packtpub.com and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Share your thoughts

Now you've finished *The Infinite Retina, Second Edition*, we'd love to hear your thoughts! If you purchased the book from Amazon, please [click here to go straight to the Amazon review page](#) for this book and share your feedback or leave a review on the site that you purchased it from.

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Index

Symbols

3D data visualization 269-271
3D Gaussian splats 176
3D data 176
3DOF headset 140
3D sensors 52, 55
5G 12, 178
6D.ai 53, 175
6DOF headset 140, 141
use cases 144
42 Silicon Valley 288, 289
360-degree cameras 171
2017 Consumer Electronics Show 297

A

adaptive AR applications 78-80
Adobe Aero 175
advanced 5G 12
Advanced Robotics for Manufacturing (ARM) 203
Agricultural Bank of China 277
AI, and data
competitive advantage, through personalized experiences 354, 355

customer behavior 353
entertainment, revolutionizing with AI-driven Spatial Computing 354
Spatial Computing 353, 354
AI chip 162
AI lingo 170
Alchera 278
Alef Aeronautics 115
Alexa 35, 51, 130, 180
Alibaba 280
Alipay 280
Amazon 56-58, 428
Amazon Alexa 61
Amazon Echo 41, 61
Amazon Go 214
Amazon Scout 228
Amazon Web Services 56
Apple 27, 53, 110, 163
glasses 43
Apple II 33
Apple Vision Pro 32, 43, 48, 168, 271, 381
business models and ecosystems 399, 400
capabilities, expanding 405
challenges and criticisms 400

content creation and development transformations 393-395

economic and market impact 396

future 411

impact 411

industry applications 388

integration, of IoT 405

integration, with AL and ML 405

mainstream adoption 412

Spatial Computing market impact 398

technological features 382

trajectory of immersive technology 411

use cases 407

work and communication, revolutionizing 412

Apple Vision Pro, for developers and innovators

- advocate, for ethical development 413
- diverse use cases 413
- visionOS Ecosystem, embracing 412

Apple Vision Pro impact, on Spatial Computing

- ecosystem impact 406
- integration, with emerging technologies 405
- Mixed Reality, advancing 404

Apple Vision Pro, industry applications

- education 391
- entertainment and media 389
- healthcare 390
- productivity and collaboration 389, 390
- retail and e-commerce 392

Apple Watch 263

Aptiv 110

AR cloud 165

- 3D Gaussian splats 176
- 3D mapping 174
- 5G 178

map 174

map building 175, 176

ARCore 152, 175

AR, from mobile to MR 150

- accommodation 159
- vergence 159
- visual displays 150

ARKit 152

artificial intelligence (AI) 13, 59-61

AR View 236

ATM security 275-280

audio augmentation 179

Augmented Reality (AR) 210, 286

- informed robots 196-203
- limitations 151-153
- origins 151-153
- visions 196
- wearable AR 154-156

Aurora 110, 111

Autodesk 205

automated content generation, immersive environments 347

- architecture and real-estate 349, 350
- automotive industry 351, 352
- education and training 350
- film and entertainment 348, 349
- game development 347
- marketing and retail 351

autonomous trucks (ATs) 111-113

autonomous vehicle (AV) companies

- Apple 110
- Aptiv 110, 111
- Aurora 110, 111
- autonomous trucks (ATs) 111-113
- Flying Cars (eVTOLs) 113-115

Ford 105
General Motors' Cruise 106
Hyundai 110, 111
Pony.ai 110, 111
Tesla 109
Waymo 107, 108
Zoox 102-105

autonomous vehicle (AV), future city
AI assistants 125, 126
auto dealer 122
entertainment 126
fewer cars 126
green thinking 125
interstates 123
navigation 125
real estate 124
remote working 124
smart roads 122, 123
traditional factories 123
travel times and ranges 124

autonomous vehicle (AV) 101, 102
cheaper car transportation costs 99-101
electronic cars, impact 115-121
future city 121, 122
future vision 93-95
paths 95-99

Autopilot 26

AWS' Amazon Sumerian 238

Azure Kinect sensor/camera 195

B

Bel Geddes, Norman 96

B&H Photo 213

biometric-multifactor identity 50

Blackberries 39

Boise 184

brainier learning 289-291

brain tricks 258

Brev.ai 344

bubbleland 339

Bubble Network 339

business models, emergence 361
digital assets 362
economic impact, on traditional industries 362
future projections 363
long-term economic trends 363
virtual economies 362

C

CaixaBank 278, 282

Calder, Dr. Calder 249

California Consumer Privacy Act (CCPA) 364

Caltech 269

camera magic 172, 173

challenges and considerations, Generative AI and Spatial Computing
environment impact and sustainability 374-378
ethical and privacy concerns 364
high costs and accessibility barriers 365
intellectual property (IP) and copyright issues 366-369
regulatory compliance and global standards 369-373
technical complexity and integration challenges 365, 366

workforce adaptation and skill gaps 378, 379

challenges and criticisms, Apple Vision Pro
 privacy and ethical Considerations 402
 technical limitations 400

ChatGPT 51

check engine light 263, 264

Chime 282

China Construction Bank 277

China Merchants Bank 277

Choreographic Interfaces 17, 18

CloudWalk 426

cobots 190

Collaborative Augmented Reality 299

computer vision 60, 169

computing architectures 56-58

confirmation bias 86

Consumer Electronics Show (CES) 23, 210

control mechanisms 48-51

Cornerstone 292-294

Cruise 24

customization and personalization, generative AI 66
 beyond gaming 72-74
 education and training 67-69
 future 76
 gaming 69-72
 interior design and architecture 66, 67
 personalized work environments 74-76
 retail and shopping 74

Cybercab 26

D

DALL-E 344

DAQRI 157, 158

DARPA 22

data clouds 13

data floods 13

data visualization 14

Dell 34

Deloitte 203

dementia 259

DigiLens 45, 163

Digital Ethics Center 428

digital twin 204-207, 356
 business opportunities 357
 congestion, reducing 357
 ethical and privacy considerations, of smart cities 359
 future 359
 infrastructure maintenance 357
 infrastructure optimization 358
 predictive maintenance 358
 public safety and emergency response, enhancing 358
 traffic flow, optimizing 357
 transforming, into intelligent urban systems 356
 urban planning and development consultancies 358
 utility management, enhancing 357

Dragonfly 280

dynamic content generation
 adaptive AR applications 78-80
 real-time data visualization 80, 81
 via generative AI 76, 77

E

EchoAR 58

economic impact, Generative AI 360
 automation and cost efficiency 360

enhanced productivity 361
streamlined operations 361

edge 58

education 336-338

Electric Sheep Robotics 15-17

electric vehicle (EV) 110

Electronic Medical Records (EMR)
systems 265

embodiment 139

emergence 133, 134

Epson 151

eVTOL vehicle types 113

experiential digital learning 286, 287

exponential learners 198

extended reality (XR) 13

EyeGym 249

eye sensors 50

F

Facebook 27, 37, 53, 163

Facebook-Cambridge Analytica
scandal 417-420

facial payment machines 275-280

festivals, VR
cinematic VR 186
interactive VR 186

finance 336

Flying Cars (eVTOLs) 113-115

fog 58

font-smoothing technology 11

Ford 105

Fuji Automatic Numerical
Control (FANUC) 200

FundamentalVR 253

Furness, Tom 45

future of education 295, 296
grabbable classroom 296-300
workplace training and simulation 301-308

G

Games Developer Conference (GDC) 125

Gaussian 176

General Data Protection Regulation
(GDPR) 364

General Motors 24

General Motors' Cruise 106

generative adversarial networks (GANs) 342

Generative AI 342
AI-driven customer interactions 343
content creation 342
creativity, expanding in entertainment and
arts 344
economic impact 360
education and training, revolutionizing 345
healthcare innovation and medical
research 346
media automation 342
product design and prototyping 343

Generative AI and Spatial Computing
challenges and considerations 364
future 379, 380

Generative AI and Spatial Computing,
convergence 346
automated content generation for
immersive environments 347
personalized customer experiences, through
customer experiences 352
real-time digital twins 356

generative AI application, Spatial Computing
accessibility 82, 83
bias 86
data quality 85, 86
data synthesis and simulation 84, 85
ethical considerations 87, 88
future 89, 90
stakeholders, educating 89

generative design 205

Glassholes 135

Goo Create 238

Google 34, 426, 428

Google Assistant 35, 51, 61, 130

Google Daydream headset
limitations 148

Google Geospatial Creator 175

Google Glass
changes 137, 138
defeated expectations 134, 135
limitations 136, 137

Google Home device 41

Google Maps 176

Google Nest Hub 61

Google Street View 423

Goo Technologies 238

GPT (Generative Pre-trained Transformer) 345

GPU (Graphics Processing Unit) 161

graphical interfaces 36-38

GreyOrange 225

H

Head Mounted Displays (HMDs) 196

Here Technologies 176

HoloLens 12, 43, 44, 159, 164, 165
features 165, 166

Hololight 67

Holoride 100

Homo sapiens 5

HR department
virtual arm for 309, 310

HTC 297
Vive 49, 146, 233

human robot 207

human social good 430-432

hybrid AI 13

hypervigilance 195

Hyundai 110, 111

I

identity and ownership 423
blurred faces, on Google Street View 423, 424

China and Facial Recognition 424, 425
issues with Facial Recognition, in US and Europe 426-430

IKEA 233

immersion 139, 140

immersive experiences, generative AI in AR and VR
behavioral training 65, 66
emotional engagement 65, 66
enhancements 63
hyper-realistic simulations 63-65

Industrial Design Thinking 198

Industrial Revolution 5.0 192

industry adoption, Vision Pro
corporate entertainments 410
education 410

- entertainment 410
- healthcare 409
- Institute for the Study of Knowledge Management in Education (ISKME) 295**
- interaction, in Spatial Computing**
 - accessibility features 387
 - cognitive load, reducing 388
 - collaboration, in shared spaces 387
 - eye-tracking 385
 - flat screen, migrating to 3D interfaces 386
 - gesture recognition 385
 - intuitive design principles 387
 - MR 386
 - natural User Interfaces 384
 - passthrough 386
 - redefining 384
 - transitions from screens to spaces 386
- User Experience (UX) Design 387
- UX Design 388
- iPhone 61, 195**
- K**
 - KiloBiteData 12
 - Kinect 164
 - Kitty Hawk 99
 - Kodak 15
 - Kopin 158, 163
 - Kraft, Dr. Daniel 253
 - Kroger 214
 - Kuehne, Marcus 100
- L**
 - Lambda School 289
 - learning futures 310
- LEGO 223
- Lenovo's glasses 45
- Leonardo 344
- LIDARs 22-24, 55, 164
- light field concept 169
- Lighthouses 146
- Limitless 49
- LinkedIn 37, 426
- Looking Glass 231
- Lumafield 196
- Lumus lenses 163
- Lyft 99
- Lyric Robot 203, 206
- M**
 - Macintosh 131, 158
 - Magic Leap 43, 159, 166, 167
 - Spatial Computing headset 14
 - manufacturing 332
 - Mapbox 52, 176
 - mapping 51-55
 - measurement bias 86
 - MediView 253, 254
 - AI integration and future directions 256, 257
 - FDA clearance and regulatory milestones 254
 - partnership with GE 255
 - surgical navigation and cardiology 256
 - XR products and technological advancements 255
 - Megvii 426
 - Merge Cube 298
 - Meta 34, 138, 141, 142, 154, 155, 426
 - Metaio 151

M

- Meta Quest 1 and 2 43
- Meta Quest 3 43, 72
- Microsoft 11, 27, 53, 56
- Microsoft Azure 58
- Microsoft Mixed Reality headsets
 - limitations 149, 150
- Midjourney 344
- Milgram, Paul 42
- Minecraft Earth 32
- Mixed Reality (MR)
 - benchmark, setting for hardware 404
 - device development 404
 - transition, between AR and VR 404
- Mobile 38, 39
- MobileEye 55
- Mojo Lens 273
- Mojo Vision 45, 273
- Moore's Law 37
- Morrow, Dennis 191
- Moven 275
- Musk, Elon 33

N

- NASA 267
- Netflix 428
- Neuralink 131
- Nokia 15, 39
- non-standalone 5G 12
- Nucleus VR 205

O

- Oakley's ski goggles 12
- Obsess
 - reference link 211
- Oculus 140, 297
- Oculus Go 140
 - inside-out tracking system 147
- Oculus Quest 32, 147-149
- Oculus Rift 149
- ODG 154
- OpenStreetMap 176
- optics and displays 42-46
- Oracle 56
- Osso VR 253
- outside-in tracking system 147
- Owlchemy Labs 149

P

- pain management 257
- paradigms 31-33
 - graphical interfaces 36-38
 - Mobile 38-40
 - personal computing 33-35
 - Spatial Computing 40, 41
- personal avatars 218
- pervasive and proactive future assistance 340
- pillow toss game 233
- Pioneer Rayz headphones 180
- Plus 297
- point-of-sale (POS) camera-equipped machines 279
- Pokémon GO 32
- Pony.ai 110, 111
- PrecisionOS 253
- predictability 315, 316
- Prime Directive 5
- PrimeSense 23

privacy and ethical considerations, Apple

Vision Pro

eye-tracking and data collection,
associated risks 402

immersive technologies 403

innovation and responsibility 403

privacy and security

Apple and FBI 422

Facebook-Cambridge Analytica
scandal 416-420

OJ Simpson and Location Data 421

Q

Qualcomm chipset 32

R

Rackspace 56

Radius Innovation 191

Ray-Ban Meta Wayfarer 138, 151

real-time data visualization 80, 81

REDCap 13

reminiscent therapy 259

Rensselaer Polytechnic Institute 299

replicating reality 6

retail change 212-215

impact 246

retail Spatial Computing experiences

future 241-246

retail theme 215-217

robot consumers 209, 210

technologies 210

three-dimensional visualization 210-212

robots 190

augmentation 190-196

role-playing games (RPGs) 69

Runway 344

Russell, Austin 53

S

sampling bias 86

Samsung HMD Odyssey 149

schizophrenia 259

Sears, Roebuck and Company 241

Second Life 281

SenseTime 426

sensors 13, 51-55, 162

Shinhan Bank 278

Siemens 198, 205

silver tsunami 193

**Simultaneous Location And
Mapping (SLAM)** 53, 230

Siri 35, 61, 130, 180

Snapchat 138

Society of Automotive Engineers (SAE) 101

Sony 45

spatial applications

anticipation for future updates 409

challenges and opportunities 408

concerns and criticisms 409

consumer reactions 408

future, building 407, 408

industry adoption 409

positive reception 408

spatial audio 180, 181

Spatial Computing 5, 35, 41, 42, 210, 269

and Vision Pro 403

benefits 7, 8

Delivery Methods 226-230

for Logistics 224-226
integrating 217
potential dangers 8-10
retail future 222-224
technologies 217-222

Spatial Computing ecosystem
growth, accelerating 407
spatial web 406

Spatial Computing Marketing Experiences 231
AR View 238-241
impact 235-238
living room 231
personalized experience 232, 233
VR World 233-235

SpinVector 198

splatting 177

standalone 5G 12

Starship Technologies 226

Suno 344

super-wide field of view (FOV) 43

Surgical Theater 253

SyncReality 218

T

Tech Mahindra 200

technical limitations, Apple Vision Pro
accessibility, for diverse demographics 402
adoption barriers 401
battery life 400
cost concerns 401
processing power and heat management 401
weight and comfort 401

technological features, Apple Vision Pro
computing paradigms 383

content creation 384
developer empowerment 383
Ecosystem Integration 384
gesture control 382
micro-OLED displays and hardware advancements 383
passthrough capabilities 382
user interface innovations 382
vision, for immersive technology 383
visionOS 382

technologies
AI 59-61
computing architectures 56-58
control mechanisms 48-51
optics and displays 42-46
sensors and mapping 51-55
wireless and communications 46, 47

Technology, Media, and Telecommunications (TMT) 127-331
2014's new wave 144-147
bandwidth, increasing 130
embodiment 139
emergence 133
evolution 131-141
Google Glass 134
immersion 139, 140
inside-out tracking 147, 148
market evolution 142-144

Tencent 280

Tesla 109, 110

Tesla car 52

Tesla Model 3 25

Tesla's Full Self-Driving (FSD) 26

Thrun, Sebastian 99

Torrealba, Eduardo 196

Transformative Experience 416

transportation system 327-330

Treos 39

Trimble 67

Tripp VR 260

Trueface 283

Twitter 37, 426

U

Uber 99

unified exponential learning system 339, 340

User Experience (UX) design 387

advancing, in Spatial Computing 388

V

Varjo 46, 138

V-cam technology 187

virtual arm

for HR department 310

Virtual Artist product 211

virtual assistants 317, 318

connected 324-326

pervasive 319, 320

proactive 321-323

virtual beings 61

virtual branch functionality

customer service 281-283

virtual coach 263, 264

virtual healthcare 247

data 265-268

MediView 254

surgery preparation and surgery 251-254

Virtualist 248-250

Virtualitics 269

virtual medicine 257-262

Virtual Reality Kitchen Experience 233

Virtual Reality

(VR) 21, 181, 182, 257-263, 286

cinematic and interactive experiences 185, 186

community building 184

content production 186

games and social experiences 183

previsualization 187

virtual trading 270-275

virtual try-ons 218

VisiCalc 35

Vision-Guided Vehicles (VGVs) 224

Visualix 54

volumetric video

shooting 170-172

Vuzix 151

W

Wayfair 233, 240

Waymo 24, 107, 108

WebAR 238

WebGL 238

WebVR 238

Wilson, Andy 49

wireless capabilities 162

wireless communications 46, 47

word processor 5

Worldwide Developers Conference (WWDC) 48

X

X 130
XR 30 255
XR 50 255
XR 90 255
XREAL 167

Y

Yitu 426
YouTube 27

Z

Zendo 266
Zoox 27, 102-105

Download a free PDF copy of this book

Thanks for purchasing this book!

Do you like to read on the go but are unable to carry your print books everywhere?

Is your eBook purchase not compatible with the device of your choice?

Don't worry, now with every Packt book you get a DRM-free PDF version of that book at no cost.

Read anywhere, any place, on any device. Search, copy, and paste code from your favorite technical books directly into your application.

The perks don't stop there, you can get exclusive access to discounts, newsletters, and great free content in your inbox daily.

Follow these simple steps to get the benefits:

1. Scan the QR code or visit the link below:



<https://packt.link/free-ebook/9781836204831>

2. Submit your proof of purchase.
3. That's it! We'll send your free PDF and other benefits to your email directly.

